arXiv:1702.02809v2 [cs.DB] 12 Apr 2017

NetClus: A Scalable Framework for Locating Top-K Sites for Placement of Trajectory-Aware Services

Shubhadip Mitra[1] Priya Saraf[1] Richa Sharma[1] Arnab Bhattacharya[1] Sayan Ranu[2]
[1]Dept. of Computer Science and Engineering, Indian Institute of Technology, Kanpur, India
[2]Dept. of Computer Science and Engineering, Indian Institute of Technology, Delhi, India
[1]{smitr,priyas,richash,arnabb}@cse.iitk.ac.in [2]sayanranu@cse.iitd.ac.in

ABSTRACT

Facility location queries identify the best locations to set up new facilities for providing service to its users. Majority of the existing works in this space assume that the user locations are *static*. Such limitations are too restrictive for planning many modern real-life services such as fuel stations, ATMs, convenience stores, cellphone base-stations, etc. that are widely accessed by mobile users. The placement of such services should, therefore, factor in the *mobility patterns* or *trajectories* of the users rather than simply their static locations. In this work, we introduce the *TOPS (Trajectory-Aware Optimal Placement of Services)* query that locates the best k sites on a road network. The aim is to optimize a wide class of objective functions defined over the user trajectories. We show that the problem is NP-hard and even the greedy heuristic with an approximation bound of $1 - 1/e$ fails to scale on urban-scale datasets. To overcome this challenge, we develop a *multi-resolution clustering* based indexing framework called NETCLUS. Empirical studies on real road network trajectory datasets show that NETCLUS offers solutions that are comparable in terms of quality with those of the greedy heuristic, while having practical response times and low memory footprints. Additionally, the NETCLUS framework can absorb dynamic updates in mobility patterns, handle constraints such as site-costs and capacity, and existing services, thereby providing an effective solution for modern urban-scale scenarios.

Keywords

Spatio-temporal databases; Facility location queries; Optimal location queries; Road networks; Trajectory-aware services;

1. INTRODUCTION AND MOTIVATION

Facility Location queries (or *Optimal location (OL) queries*) in a road network aim to identify the best locations to set up new facilities with respect to a given service [11, 14, 17, 29, 36, 40]. Examples include setting up new retail stores, gas stations, or cellphone base stations. OL queries also find applications in various spatial decision support systems, resource planning and infrastructure management [25, 35].

With growing applications of data-driven location-based systems, the importance of OL queries is well-recognized in the database community [11, 40, 45]. However, most of the existing works assume the users to be *fixed* or *static*. Such an assumption is often too prohibitive. For example, services such as gas stations, ATMs, bill boards, traffic monitoring systems, etc. are widely accessed by users while commuting. Further, it is common for many users to make their daily purchases while returning from their offices. Consequently, the placement of facilities for these services require taking into consideration the mobility patterns (or trajectories) of the users rather than their static locations. We refer to such services as *trajectory-aware services*. Formally, a *trajectory* is a sequence of location-time coordinates that lie on the path of a moving user. Note that trajectories strictly generalize the static users' scenario because static users can always be modeled as trajectories with a single user location. Thus, trajectories capture user patterns more effectively. Such trajectory data are commonly available from GPS traces [20] or CDR (Call Detail Records) data [25] recorded through cellphones, social network check-ins, etc. Recently, there have been many works in the area of trajectory data analytics [9, 16, 21–23, 27, 30, 38, 39].

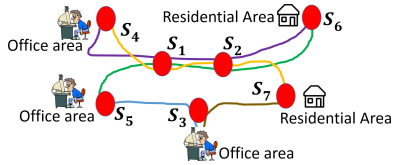

Figure 1: Illustration of the need of trajectory-aware querying for optimal locations. Each line represents a separate user trajectory.

To illustrate the need for trajectory-aware optimal location queries, consider Fig. 1. There are 5 locations that are either homes or offices and 5 trajectories of users commuting across them. A company wants to open 2 new gas stations. For simplicity, we assume that a trajectory of a user is *satisfied* if it passes through at least one gas station. If only the static locations are considered, i.e., any two out of the five office and residential areas are to be selected, no combination of gas stations would satisfy all the users. In contrast, if we factor in the mobility of the users, choosing S_1 and S_3 as the installation locations satisfies all the trajectories of the users.

Note that it is not enough to simply look at trajectory counts in each possible installation location and then choose the two most frequent ones (S_1 and S_2). The combination may not be effective since a large number of trajectories between them can be common, thereby reducing each other's utilities.

In this work, we formalize the problem of OL queries over trajectories of users in road networks. We refer to this as *TOPS (Trajectory-aware Optimal Placement of Services)* query. Given a set of user trajectories $\mathcal{T}$, and a set of candidate sites $\mathcal{S}$ over a road network that can host the services, TOPS query with input parameters k and preference function ψ seeks to report the best k sites $\mathcal{Q} \subseteq \mathcal{S}$ maximizing a *utility function* that is defined over the preference function ψ which captures how a particular candidate site is *preferred* by a

given trajectory.

Facility location queries with respect to trajectories have been studied by a number of previous works [2–5, 7, 28]. Although the formulations are not identical, the common eventual goal is to identify the best sites. However, all these works remain limited to a theoretical exercise and cannot be applied in a real-life scenario due to a number of issues as explained next.

• **Data-based mobility model:** Existing techniques are neither based on real trajectories nor on real road networks [2–5, 7]. They base their solutions on simplistic assumptions such as traveling in shortest paths and synthetic road networks. It is well known that the shortest path assumption does not hold in real life [1]. In many works, the distances are *not* computed over the road network but approximated using some spatial distance measure such as L_2 norm. Our framework is the first to study TOPS on real trajectories over real road networks.

• **Generic framework:** We develop the first generic framework to answer TOPS queries across a wide family of preference functions ψ that are non-increasing w.r.t. the distance between any pair of trajectory and candidate site. The proposed framework encompasses many of the existing formulations and also considers other practical factors such as capacity constraints, site-costs, dynamic updates, etc.

• **Scalability:** The state-of-the-art technique for a basic version of TOPS query [2] requires prohibitively large memory. Consequently, it fails to scale on urban-scale datasets (Further details in Sec. 8). Hence, a scalable framework for TOPS query is a basic necessity. In addition, all OL queries including TOPS are typically used in an interactive fashion by varying the various parameters such as k and the coverage threshold τ [11]. Moreover, in certain ventures, such as deployment of mobile ATM vans(https://goo.gl/WjSPvx), real-time answers are need based on current trajectory patterns. Hence, practical response time with the ability to absorb data updates is critical. This factor has been completely ignored in the existing works.

• **Extensive benchmarking:** Since TOPS queries and their variants are NP-hard, heuristics have been proposed. How do their effectiveness vary across road-network topologies? Are these heuristics biased towards certain specific parameter settings? The existing techniques are generally silent on these questions. We, on the other hand, perform benchmarking that is grounded to reality by extensively studying the performance of TOPS across multiple major city topologies and other important parameters.

To summarize, the proposed framework is the first practical and generic solution to address TOPS queries. Fig. 2 depicts the top-level flow diagram of our solution. Given the raw GPS traces of user movements, they are map-matched [33] to the corresponding road network. Using the map-matched trajectories, a multi-resolution clustering of the road network is built to construct the index structure NETCLUS. Indexed views of both the candidate sites and the trajectories are maintained in a compressed format at various granularities. This completes the offline phase. In the online phase, given the query parameters, the optimum clustering resolution to answer the query is identified, and the corresponding views of the trajectories and road network are analyzed to retrieve the best k sites for facility locations.

The major contributions of our work are as follows:

1. We propose a highly generic trajectory-aware optimal placement of services problem, TOPS, that can handle the wide class of preference functions ψ that are non-increasing w.r.t. the distance between any pair of trajectory and candidate site (Sec. 2).
2. TOPS is NP-hard, and even the greedy approach does not scale in city-scale datasets (Sec. 3). To overcome this bottleneck, we design a *multi-resolution clustering-based* index structure called NETCLUS that generates a *compressed* representation of the road network and the set of trajectories (Sec. 4 and Sec. 5). The solutions returned by NETCLUS framework have bounded quality guarantees.
3. The proposed TOPS formulation is capable of absorbing dynamic updates (Sec. 6) and manage realistic constraints such as cost, capacity, and existing services. (Sec. 7).
4. Extensive experiments on real datasets show that NETCLUS offers solutions that are comparable in terms of quality with those of the greedy heuristic, while having practical response times and fairly low memory footprints (Sec. 8).

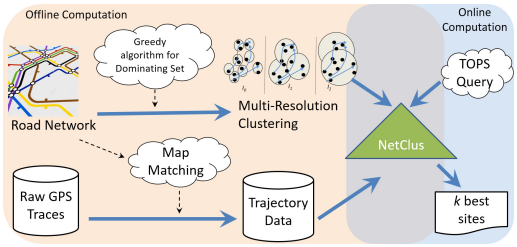

Figure 2: Flowchart of TOPS querying framework.

Symbol	Description
$G = (V, E), \lvert V\rvert = N$	Road network G with node set V and edge set E
$\mathcal{T}, \lvert\mathcal{T}\rvert = m$	Set of trajectories
$\mathcal{S}, \lvert\mathcal{S}\rvert = n$	Set of candidate sites
$d(u, v)$	Distance of shortest path from node u to v
$d_r(u, v)$	Round-trip distance between nodes u and v
$d_r(T_j, s_i)$	Round-trip distance from trajectory T_j to site s_i
k	Desired number of service locations
τ	Coverage threshold
$\psi(T_j, s_i)$	Preference function for trajectory T_j and site s_i
$\mathcal{Q} \subseteq \mathcal{S}, \lvert\mathcal{Q}\rvert = k$	Set of locations to set up service
U_j	Utility of trajectory T_j over the set of sites $\mathcal{Q}$
$U(\mathcal{Q}) = \sum_{j=1}^{m} U_j$	Total utility offered by $\mathcal{Q}$

Table 1: Symbols used in the paper.

2. THE TOPS PROBLEM

Consider a road network $G = \{V, E\}$ over a geographical area where V denotes the set of road intersections, and E denotes the set of road segments between two adjacent road intersections. The direction of the underlying traffic that passes over a road segment is modeled by directed edges.

Assume a set of candidate sites $\mathcal{S} = \{s_1, \cdots, s_n\} \subseteq V$ where a certain service or facility can be set up. The choice of $\mathcal{S}$ is generally provided by the application itself by taking into account various factors such as availability, reputation of neighborhood, price of land, etc. Most of these factors are outside the purview of the main focus of this paper and are, therefore, not studied. We simply assume that the set $\mathcal{S}$ is provided as input to our problem. However, as described later, if all the latent factors of choosing a site can be encapsulated as its cost, we can handle it quite robustly.

The candidate sites can be located anywhere on the road network. If it is already on a road intersection, then it is part of the set of vertices V. If not, i.e., if it is on the middle of a road connecting vertices u and v, without loss of generality, we augment V to include this site as a new vertex w. We augment the edge set E by two new edges (u, w) and (v, w) (with appropriate directions) and remove the old edge (u, v). Thus, ultimately, $\mathcal{S} \subseteq V$.

The set of candidate sites $\mathcal{S}$ can be in addition to the set of existing service locations $\mathcal{E}_{\mathcal{S}}$.

The set of trajectories over the road network is denoted by $\mathcal{T} = \{T_1, \cdots, T_m\}$ where each trajectory is a sequence of nodes, $T_j = \{v_{j_1}, \cdots, v_{j_l}\}$, $v_{j_i} \in V$. The trajectories are usually recorded as GPS traces and may contain arbitrary spatial points on the road network. For our purpose, each trajectory is map-matched [33] to form a sequence of road intersections through which it passes. We assume that each trajectory belongs to a separate user. However, the framework can easily generalize to multiple trajectories belonging to a single user by taking union of each of these trajectories.

Trajectories	Preference scores for different sites		
	s_1	s_2	s_3
T_1	$\psi(T_1, s_1) = 0.4$	$\psi(T_1, s_2) = 0.11$	$\psi(T_1, s_3) = 0$
T_2	$\psi(T_2, s_1) = 0$	$\psi(T_2, s_2) = 0.5$	$\psi(T_2, s_3) = 0.6$

Table 2: Examples of trajectories with site preferences.

Suppose $d(v_i, v_j)$ denotes the shortest network distance along a directed path from node v_i to v_j, and $d_r(v_i, v_j)$ denotes the shortest distance of a round-trip starting at node v_i, visiting v_j, and finally returning to v_i, i.e., $d_r(v_i, v_j) = d(v_i, v_j) + d(v_j, v_i)$. In general, $d(v_i, v_j) \neq d(v_j, v_i)$, but $d_r(v_i, v_j) = d_r(v_j, v_i)$. With a slight abuse of notation, assume that $d_r(T_j, s_i)$ denotes the *extra* distance traveled by the user on trajectory T_j to avail a service at site s_i. Formally, $d_r(T_j, s_i) = \min_{\forall v_k, v_l \in T_j} \{d(v_k, s_i) + d(s_i, v_l) - d(v_k, v_l)\}$.

It is convenient for a user to avail a service only if its location is not too far off from her trajectory. Thus, beyond a distance τ, we assume that the utility offered by a site s_i to a trajectory T_j is 0. We call this user-specified distance τ as the *coverage threshold*.

DEFINITION 1 (COVERAGE). *A candidate site s_i* covers *a trajectory T_j if the distance $d_r(T_j, s_i)$ is at most τ, where $\tau \geq 0$ is the* coverage threshold.

For all sites within the coverage threshold τ, the user also specifies a preference function ψ. The preference function $\psi(T_j, s_i)$ assigns a score (normalized to $[0, 1]$) for a trajectory T_j and a site s_i that indicates how much s_i is preferred by the user on trajectory T_j. Higher values indicate higher preferences with 0 indicating no preference. In general, sites that are *closer* to the trajectory have *higher* preferences than those farther away. Usage of such preference functions are common in location analysis literature [44].

DEFINITION 2 (PREFERENCE FUNCTION ψ). *$\psi : (\mathcal{T}, \mathcal{S}) \rightarrow [0, 1]$ is a real-valued preference function defined as follows:*

$$\psi(T_j, s_i) = \begin{cases} f(d_r(T_j, s_i)) & \text{if } d_r(T_j, s_i) \leq \tau \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where f is a non-increasing function of $d_r(T_j, s_i)$.

The goal of TOPS query is to report a set of k sites $\mathcal{Q} \subseteq \mathcal{S}$, $|\mathcal{Q}| = k$, that maximizes the preference score over the set of trajectories. The preference score of a trajectory T_j over a set of sites $\mathcal{Q}$ is defined as the *utility function* U_j for T_j, which is simply the *maximum* score corresponding to the sites in $\mathcal{Q}$, i.e., $U_j = \max_{s_i \in \mathcal{Q}} \{\psi(T_j, s_i)\}$. The various symbols used in the TOPS formulation are listed in Table 1.

The generic TOPS query formulation is stated next.

PROBLEM 1 (TOPS). *Given a set of trajectories $\mathcal{T}$, a set of candidate sites $\mathcal{S}$ that can host the services, TOPS problem with query parameters (k, τ, ψ) seeks to report the best k sites, $\mathcal{Q} \subseteq \mathcal{S}$, $|\mathcal{Q}| = k$, that maximize the sum of trajectory utilities, i.e., $\mathcal{Q} = \arg\max \sum_{j=1}^{m} U_j$ where $U_j = \max_{s_i \in Q} \{\psi(T_j, s_i)\}$.*

Extensions and variants of TOPS: The TOPS framework as defined above is *generic*. It can have various extensions as described in Sec. 7.1 and Sec. 7.2 and can work with existing facilities as well (Sec. 7.3). Finally, the preference function in Def. 2 subsumes various existing models as exemplified in Sec. 7.4.

2.1 Properties of TOPS

TOPS problem is NP-hard. To show that, we first define a specific instance of TOPS where the preference function $\psi(T_j, s_i)$ is a binary function. A binary function is a natural choice in situations where a service provider wants to intercept the maximum number of trajectories that pass within a vicinity of it [2].

Algorithm	Selected sites $\mathcal{Q}$	Utility U
OPTIMAL	$\{s_1, s_3\}$	1.0
INC-GREEDY	$\{s_1, s_2\}$	0.9

Table 3: Utilities of different algorithms for Example 1 with $k = 2$.

DEFINITION 3 (BINARY INSTANCE OF TOPS). *TOPS with parameters (k, τ, ψ) where the preference score is given by:*

$$\psi(T_j, s_i) = \begin{cases} 1 & \text{if } d_r(T_j, s_i) \leq \tau \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

THEOREM 1. *TOPS is NP-hard.*

PROOF. This follows from the fact that the binary instance of TOPS is NP-hard, as set cover problem is reducible to it [2]. □

Next we show that the sum of utilities $U = \sum_{j=1}^{m} U_j$ is a *non-decreasing sub-modular* function. A function f defined on any subset of a set $\mathcal{S}$ is *sub-modular* if for any pair of subsets $\mathcal{Q}, \mathcal{R} \subseteq \mathcal{S}$, $f(\mathcal{Q}) + f(\mathcal{R}) \geq f(\mathcal{Q} \cup \mathcal{R}) + f(\mathcal{Q} \cap \mathcal{R})$ [37].

THEOREM 2. *$U = \sum_{j=1}^{m} U_j$ is non-decreasing and sub-modular.*

PROOF. Consider any two pair of subsets $\mathcal{Q}, \mathcal{R} \subseteq \mathcal{S}$ such that $\mathcal{Q} \subseteq \mathcal{R}$. Since $U_j = \max_{s_i \in \mathcal{Q}} \{\psi(T_j, s_i)\}$, $U_j(\mathcal{R}) \geq U_j(\mathcal{Q})$. Therefore, $\sum_{j=1}^{m} U_j(\mathcal{R}) = U(\mathcal{R}) \geq \sum_{j=1}^{m} U_j(\mathcal{Q}) = U(\mathcal{Q})$. Hence, U is a non-decreasing function.

To show that U is sub-modular, following [37], it is sufficient to show that for any two subsets $\mathcal{Q}, \mathcal{R} \subseteq \mathcal{S}$, $\mathcal{Q} \subset \mathcal{R}$, for any site $s \in \mathcal{S} \setminus \mathcal{R}$, $U(\mathcal{Q} \cup \{s\}) - U(\mathcal{Q}) \geq U(\mathcal{R} \cup \{s\}) - U(\mathcal{R})$. For our formulation, it is, thus, enough if for any trajectory $T_j \in \mathcal{T}$,

$$U_j(\mathcal{Q} \cup \{s\}) - U_j(\mathcal{Q}) \geq U_j(\mathcal{R} \cup \{s\}) - U_j(\mathcal{R}) \quad (3)$$

Suppose the site $s^* \in \mathcal{R} \cup \{s\}$ offers the maximal utility to trajectory T_j in the set of sites $\mathcal{R} \cup \{s\}$. There can be two cases:
(a) $s^* = s$: $U_j(\mathcal{Q} \cup \{s\}) = U_j(\{s\}) = U_j(\mathcal{R} \cup \{s\})$. Further, $U_j(\mathcal{Q}) \leq U_j(\mathcal{R})$ since $U_j = \{\max\{\psi(T_j, s_i)\} | s_i \in Q\}$. These two inequalities together lead to Ineq. (3).
(b) $s^* \neq s$: Here, $U_j(\mathcal{R} \cup \{s\}) - U_j(\mathcal{R}) = 0$. As $U_j()$ is non-decreasing, $U_j(\mathcal{Q} \cup \{s\}) \geq U_j(\mathcal{Q})$. Thus, Ineq. (3) follows. □

3. ALGORITHMS FOR TOPS

We first state an example, using which we will explain the different algorithms for TOPS problem.

EXAMPLE 1. Consider Table 2, which lists 2 trajectories T_1, T_2 with their preference scores for 3 sites s_1, s_2 and s_3. We set $k = 2$ sites to be selected for TOPS query. The utilities achieved by different algorithms (discussed next) are shown in Table 3. The optimal solution consists of the sites s_1 and s_3 with a total utility of 1.0.

3.1 Optimal Algorithm

We first discuss the *optimal* solution to TOPS problem in the form of an *integer programming problem* (ILP):

$$\text{Maximize } U = \sum_{j=1}^{m} U_j \quad (4)$$

$$\text{such that } \sum_{i=1}^{n} x_i \leq k, \quad (5)$$

$$\forall j = 1, \ldots, m,\ U_j \leq \max_{1 \leq i \leq n} \{\psi(T_j, s_i) \times x_i\} \quad (6)$$

$$\forall i = 1, \ldots, n,\ x_i \in \{0, 1\} \quad (7)$$

The above IP is, however, not in the form of a standard integer linear program (ILP). To do so, the constraints in Ineq. (6) can be linearized, as discussed in Appendix A.1. Since TOPS is NP Hard,

the cost of solving this optimal algorithm is *exponential* with respect to the number of trajectories m, and the number of sites n and, therefore, *impractical* except for very small m and n. This is demonstrated through experiments in Sec. 8.3.

3.2 Approximation Algorithms

We next present a greedy approximation algorithm. Before that, let us define a few terms and how they are computed.

Given a site $s_i \in \mathcal{S}$, $TC(s_i)$ denotes the set of trajectories covered by s_i, i.e., $TC(s_i) = \{T_j | d_r(T_j, s_i) \leq \tau\}$. Similarly, given a trajectory T_j, $SC(T_j)$ denotes the set of sites that cover T_j, i.e., $SC(T_j) = \{s_i | d_r(T_j, s_i) \leq \tau\}$. For each site s_i, the site weight $w_i = \sum_{j=1}^{m} \{\psi(T_j, s_i) | T_j \in TC(s_i)\}$ denotes the sum of preference scores of all the trajectories covered by it. Table 4 tabulates some of the important notations.

Since τ is available only at query time, the computation of these covering sets need to be efficient. Hence, for each site $s_i \in \mathcal{S}$, the shortest path distances to all other nodes in V are pre-computed. This requires $O(|\mathcal{S}|.|E| \log |V|)$ time [12]. As the road network graphs are almost planar, i.e., $|E| = O(|V|)$, the above cost simplifies to $O(n.N \log N)$ where $n = |\mathcal{S}|$ and $N = |V|$.

Using these distance values, the distance between each pair of site and trajectory, $d_r(T_j, s_i)$, is computed. Suppose there are at most l nodes in a trajectory T_j, i.e., $\max_{T_j \in \mathcal{T}} |T_j| = l$. Referring to the definition of $d_r(T_j, s_i)$ in Sec. 2, it follows that $d_r(T_j, s_i)$ can be computed in $O(l^2)$ time for a given site-trajectory pair using the above site-node distances. Therefore, the total cost across all n sites and m trajectories is $O(mnl^2)$.

Each site s_i maintains the set of trajectories in an ascending order of $d_r(T_j, s_i)$ while each trajectory T_j maintains the sites in a similar ascending order based on $d_r(T_j, s_i)$ again. Sorting increases the time to $O(mn(\log n + \log m))$. The space complexity for storing all site to trajectory distances is $O(mn)$.

At query time, when τ is available, the sets TC, SC, and the site weights can be, thus, computed from the stored distance matrix in $O(mn)$ time.

Alternatively, the sets TC and SC can be computed at query time when τ is made available. However, for a city-sized road network, even a small τ includes a lot of neighbors (e.g., for the Beijing dataset described in Sec. 8, there are more than 40 sites on average within a distance of $\tau = 0.8$ Km from a site). Since this needs to be computed for all the candidate sites, this approach is infeasible at query time.

3.3 Inc-Greedy

We adapt the general greedy heuristic for maximizing non-decreasing sub-modular functions [37] to design INC-GREEDY. The main idea is based on the principle of *maximizing marginal gain*.

Then the algorithm starts with an empty set of sites $\mathcal{Q}_0 = \varnothing$, and incrementally adds the sites (one in each of the k iterations) such that each successive addition of a site produces the maximal marginal gain in the utility U. More specifically, if the set of selected sites after $\theta - 1$ iterations ($\theta = 1, \ldots, k$) is $\mathcal{Q}_{\theta-1} = \{s_1, \ldots, s_{\theta-1}\}$, in the θ^{th} iteration, INC-GREEDY chooses the site $s_\theta \notin \mathcal{Q}_{\theta-1}$ such that $U(\mathcal{Q}_\theta = \mathcal{Q}_{\theta-1} \cup \{s_\theta\}) - U(\mathcal{Q}_{\theta-1})$ is *maximal*.

The marginal utility gained due to the addition of s_i to the set $\mathcal{Q}_{\theta-1}$ in iteration $\theta = 1, \ldots, k$ is $U_\theta(s_i) = U(\mathcal{Q}_{\theta-1} \cup \{s_i\}) - U(\mathcal{Q}_{\theta-1})$. Thus, $U(\mathcal{Q}_\theta) = \sum_{i=1}^{\theta} U_i(s_i)$. If α_{ji} denotes the marginal gain in the utility of trajectory T_j by addition of site s_i to the existing set of chosen sites, then, at the end of each iteration $\theta = 1, \ldots, k$, $U_\theta(s_i) = \sum_{\forall T_j \in TC(s_i)} \alpha_{ji}$.

INC-GREEDY algorithm begins by computing the sets TC, SC, and the site-weights as explained above. It then initializes the marginal utilities of each site $s_i \in \mathcal{S}$ to its site weight, i.e., $U_1(s_i) = w_i$. Further, for each trajectory $T_j \in TC(s_i)$, it initializes $\alpha_{ji} = \psi(T_j, s_i)$.

In iteration $\theta = 1, \ldots, k$, it selects the site s_θ with *maximal* marginal utility, i.e., $\forall s_i \notin \mathcal{Q}_{\theta-1}, U_\theta(s_\theta) \geq U_\theta(s_i)$. If multiple candidate sites yield the same maximal marginal gain, it chooses the one with maximal weight. Still, if there are ties, then without loss of generality, the site with the highest index is selected.

For each trajectory $T_j \in TC(s_\theta)$, first the utility U_j is updated as follows: $U_j \leftarrow \max(U_j, \psi(T_j, s_\theta))$. In case there is a gain in U_j, the marginal utility of each site $s_i \in SC(T_j)$ is updated.

The pseudocode of INC-GREEDY is prsented in Appendix A.2.

In Example 1, at iteration 1, the site s_2 has the largest marginal utility of 0.61 and is, therefore, selected. With $\mathcal{Q}_1 = \{s_2\}$, the marginal utility of adding the site s_1 is $0.4 - 0.11 = 0.29$, and that of site s_3 is $0.6 - 0.5 = 0.1$. Hence, INC-GREEDY selects the sites s_1 and s_2, and yields an utility $U = 0.9$, as indicated in Table 3.

We next analyze the quality of INC-GREEDY.

LEMMA 1. *The approx. bound of* INC-GREEDY *is* $(1 - 1/e)$.

PROOF. Since INC-GREEDY is a direct adaptation of the generic greedy algorithm for maximizing non-decreasing sub-modular functions under cardinality constraints [37] with $U(\varnothing) = 0$, it offers the approximation bound of $\left(1 - (1 - \frac{1}{k})^k\right) \geq \left(1 - \frac{1}{e}\right)$. □

LEMMA 2. $U(\mathcal{Q}_k) \geq (k/n).U(\mathcal{S})$.

PROOF. From Theorem 2 it follows that successive marginal utilities are non-increasing, i.e., $U_\theta(s_\theta) \geq U_{\theta+1}(s_{\theta+1})$. Thus,

$$U_\theta(s_\theta) \geq \frac{\sum_{i=\theta}^{n} U_i(s_i)}{n - \theta + 1} \geq \frac{U(\mathcal{Q}_n) - U(\mathcal{Q}_{\theta-1})}{n - \theta + 1} \tag{8}$$

Next, we claim that $\forall \theta = 1, 2, \ldots, n$,

$$U(\mathcal{Q}_\theta) \geq (\theta/n).U(\mathcal{Q}_n)$$

We will prove this by induction on θ.

Consider the base case $\theta = 1$. Using Ineq. (8), $U(\mathcal{Q}_1) = U_1(s_1) \geq U(\mathcal{Q}_n)/n$ since $U(\mathcal{Q}_0) = 0$. Thus, the induction hypothesis is true for the base case.

Next, we assume it to be true for iteration θ. For iteration $\theta + 1$,

$$\begin{aligned} U(\mathcal{Q}_{\theta+1}) &= U(\mathcal{Q}_\theta) + U_{\theta+1}(s_{\theta+1}) \\ &\geq U(\mathcal{Q}_\theta) + \frac{U(\mathcal{Q}_n) - U(\mathcal{Q}_\theta)}{n - \theta} \\ &\geq \left(\frac{n - \theta - 1}{n - \theta}\right) U(\mathcal{Q}_\theta) + \left(\frac{1}{n - \theta}\right) U(\mathcal{Q}_n) \end{aligned}$$

Using the induction hypothesis for iteration θ,

$$\begin{aligned} U(\mathcal{Q}_{\theta+1}) &\geq \left(\frac{n - \theta - 1}{n - \theta}\right) \left(\frac{\theta}{n}\right) U(\mathcal{Q}_n) + \left(\frac{1}{n - \theta}\right) U(\mathcal{Q}_n) \\ &\geq \left(\frac{\theta + 1}{n}\right) .U(\mathcal{Q}_n) \end{aligned}$$

Thus, the induction hypothesis holds true for any $\theta = 1, \ldots, n$. Since $\mathcal{Q}_n = \mathcal{S}$, after k iterations, $U(\mathcal{Q}_k) \geq (k/n).U(\mathcal{S})$. □

LEMMA 3. *The approximation bound of* INC-GREEDY *is* (k/n).

PROOF. Since U is non-decreasing (from Th. 2) and $OPT \subseteq \mathcal{S}$, therefore $U(OPT) \leq U(\mathcal{S})$. The rest follows from Lem. 2. □

THEOREM 3. *The approximation bound of* INC-GREEDY *for TOPS is* $\max\{1 - 1/e, k/n\}$.

PROOF. The result follows from Lem. 1 and Lem. 3. □

The next result establishes the complexity of INC-GREEDY.

THEOREM 4. *The time and space complexity bounds for* INC-GREEDY *are* $O(k.mn)$ *and* $O(mn)$ *respectively.*

PROOF. As discussed earlier in Sec. 3.2, the computation of the sets TC, SC and the weights require $O(mn)$ time and storage.

Initializing the marginal utilities $U_\theta(s_i)$ of all the sites and α_{ji} values for all trajectory site pairs require $O(mn)$ time.

In any iteration $\theta = 1, \ldots, k$, selecting the site with maximal marginal utility requires at most $O(n)$ time. The largest size of any TC set is at most m and that of any SC set is at most n. Hence, scanning the set of trajectories $TC(s_\theta)$ and updating the utilities requires $O(m)$ time. If U_j gets updated, then scanning the set of sites in $SC(T_j)$ and updating the values of $U_\theta(s_i)$ and α_{ji} requires $O(n)$ time. Thus, the time complexity of any iteration θ is $O(mn)$. Hence, the total time complexity of k iterations is $O(kmn)$.

Since no other space is required than that for storing TC, SC and site weights, the total space complexity is $O(mn)$. □

3.4 Limitations of Inc-Greedy

Although INC-GREEDY provides a constant factor approximation, it is not scalable to large real-life datasets. The reasons are:

High query cost: The input parameters for TOPS query, (k, τ, ψ), are available only at query time. Hence, the covering sets TC, SC, and the site weights (that depend on the value of τ) can be generated *only* at query time. Even if all pairwise site-to-trajectory distances are pre-computed, this step requires a high computation cost of $O(mn)$ (both in terms of time and memory) where m, n denote the number of trajectories, and candidate sites, respectively. Due to this reason, for real city-scale datasets (such as the Beijing dataset [42] used in our experiments that has over 120,000 trajectories and 250,000 sites), INC-GREEDY is not scalable. Fig. 6b shows that INC-GREEDY takes about 2000 sec. to complete for $\tau = 1.2$ Km. and $k = 5$, and goes out of memory for $\tau > 1.2$ Km.

High storage cost: As discussed above, to facilitate faster computation of covering sets TC, we need to pre-compute all pairwise site-to-trajectory distances. However, for any city-scale dataset, this storage requirement is prohibitively large. For example, the Beijing dataset [42] require close to 250 GB of storage. This is unlikely to fit in the main memory and, therefore, multiple expensive random disk seeks are required at run-time. Even with pre-computed distances up to 10 Km., INC-GREEDY crashes beyond τ 1.2 Km. (shown in Tab. 9.)

High update cost: INC-GREEDY is also not amenable to updates in trajectories and sites. If a new trajectory is added, its distance to all the sites needs to be computed and sorted. In addition, the sorted TC set of all the sites need to be updated as well. Similarly, if a new candidate site is added, the distances of all the trajectories to this site will need to be computed and sorted. Such costly update operations are impractical, especially at run-time.

A careful analysis of INC-GREEDY reveals that there are two main stages of the algorithm. In the first, the sets TC, SC and the site weights are computed. In the second stage, some of these sets are updated in an iterative manner.

The first stage is heavier in terms of time and space requirements. Thus, to make it efficient, we use an index structure, NETCLUS, which is described in Sec. 4 and Sec. 5. The use of indexing reduces the computational burden of the update (i.e., the second) stage as well. Further, Section 6 shows how NETCLUS allows easier handling of additions and deletions of candidate sites and trajectories.

In addition, if the preference function ψ is binary as defined in Def. 3, the update steps for INC-GREEDY can be performed quite efficiently using FM sketches [15,32]. We next describe the details.

3.5 Using FM Sketch to speed up Inc-Greedy

The main use of FM sketches is in counting the number of *distinct* elements in a set or union of sets [15]. Suppose, the maximum number of distinct elements is N. An FM sketch is a bit vector, which is initially empty, and is of size at least $O(\log_2 N)$. The probability of an element from the domain hashing into the i^{th} bit of the FM sketch is 2^{-i}. Thus, if a set has Γ elements, the probability of the last bit marked in the FM sketch is $O(\log_2 \Gamma)$. Hence, after the elements of a set are hashed to the FM sketch, the last set bit can be used to estimate the number of distinct elements in the set. (The details of how the hashing function is chosen and the exact estimates are in [15].) Although the FM sketch does not count the number of distinct elements exactly, it provides a multiplicative guarantee on the error in counting. When more copies of the FM sketch is used, the error decreases.

The FM sketch can be used to speed up the update stage of INC-GREEDY, since selecting a site with the largest maximal utility is the same as selecting a site that covers the largest number of *distinct* trajectories not yet covered.

For each site $s_i \in \mathcal{S}$, the set of trajectories that it covers, i.e., $TC(s_i)$ is maintained as an FM sketch. Thus, instead of maintaining $O(m)$-sized lists for each site where m is the total number of trajectories, we only need to maintain $O(\log_2 m)$-sized bit vectors per site.

Suppose the count of distinct trajectories covered by a site s_i is χ_i. The *marginal* utility of site s_j when site s_i has been chosen is the number of distinct trajectories that the two sites *together* cover over the number of trajectories that site s_i *alone* covers. The estimate of the number of trajectories covered by the union of s_j and s_i can be obtained by the bitwise OR of the FM sketches corresponding to s_j and s_i. If this estimate is χ_{ij}, the marginal utility of site s_j over site s_i is $\chi_{ij} - \chi_i$.

Therefore, when there are n candidate sites, to determine the site that provides the best marginal utility over site s_i, $n - 1$ such bitwise OR operations are performed, and the maximum is chosen. At the end of the θ^{th} iteration, the combined number of trajectories covered by the sites in $\mathcal{Q}_\theta$ is stored by the *union* of the FM sketches obtained successively in the θ iterations. The $(\theta+1)^{\text{th}}$ site is chosen by using this combined FM sketch as the base.

The above brute-force algorithm can be improved in the following way. The upper bound of the marginal utility for any site s_j is its own utility. Thus, if the current best marginal utility of another site s_k is already greater than that, it is not required to do the union operation with s_j. If the sites are *sorted* according to their utilities, the scan can stop as soon as the first such site s_j is encountered. All sites having a lower utility are guaranteed to be useless as well.

In our implementation, the FM sketches are stored as 32-bit words. This allows handling of roughly 2^{32} (which is more than 4 billion) number of trajectories. The length 32 is chosen since the bitwise OR operation of two such regular-sized words is extremely fast in modern operating systems.

4. OFFLINE CONSTRUCTION OF NETCLUS

As discussed in Sec. 3.4, INC-GREEDY has two computationally expensive components. While FM sketch expedites the information update component, computation of TC, SC, etc. still remain a bottleneck with $O(mn)$ time and storage complexity. To overcome this scalability issue, we develop an index structure.

One of the most natural ways to achieve the above objectives is to cluster the sites in the road network to reduce the number, and then apply INC-GREEDY on the cluster representatives. The clustering of sites can be done according to two broad strategies. The first is to cluster the sites that are highly similar in terms of their trajectory covering sets TC. The similarity between these covering sets can be quantified using the *Jaccard similarity* measure (Appendix B.1). However, this approach is not practical due to two major limitations: (1) Since the coverage threshold τ is available *only* at query time, the covering sets TC can be computed only at query time.

Hence, the clustering can be performed only at query time. This leads to impractical query time. (2) Alternatively, multi-resolution clustering may be performed based on few fixed values of τ, and at query time, a clustering instance of a particular resolution is chosen based on the value of the query parameter τ. However, this still requires computing the similarity between each pair of sites. Owing to large number of sites, and large size of the covering sets, such computation demands impractical memory overhead. Hence, we adopt the second clustering option, that of distance-based clustering.

We first state one basic observation. If two sites are close, the sets of trajectories they cover are likely to have a high overlap. Hence, when $k \ll n$, which is typically the case, the sites chosen in the answer set are likely to be distant from each other. The index structure, NETCLUS, is designed based on the above observation.

Our method follows two main phases: offline and online. In the offline phase, clusters are built at multiple resolutions. This forms the different index instances. A particular index instance is useful for a particular range of query coverage thresholds. In the online phase, when the query parameters are known, first the appropriate index instance is chosen. Then the INC-GREEDY algorithm is run with the cluster representatives of that instance.

We explain the offline phase in this section and the online phase in the next. The important notations used in the NETCLUS scheme are listed in Table 4.

4.1 Distance-Based Clustering

The clustering method is parameterized by a distance threshold R, which is the *maximum cluster radius*. The round-trip distance from any node within the cluster to the cluster-center is constrained to be at most $2R$. The radius is varied to obtain clusters at multiple resolutions. We describe the significance of the choices of R later.

The objective of the clustering algorithm is to partition the set of nodes in the road network, V, into disjoint clusters such that the number of such clusters is *minimal*. This leads not only to savings in index storage, but more importantly, it results in faster query time, as INC-GREEDY is run on a smaller number of cluster representatives. We next describe how to achieve this objective.

4.1.1 Generalized Dominating Set Problem (GDSP)

Given an undirected graph $G = (V, E)$, the *dominating set problem (DSP)* [19] computes a set $D \subseteq V$ of *minimal* cardinality, such that for each vertex $v \in V - D$, there exists a vertex $u \in D$, such that $(u, v) \in E$. DSP is NP-hard [19]. In [10], it was generalized to the *measured dominating set* problem for weighted graphs. In this work, we propose a *generalized dominating set problem (GDSP)*, that uses a different notion of dominance from [10].

PROBLEM 2 (GDSP). *Given a weighted directed graph $G = (V, E, W)$ where $W : E \to \mathbb{N}$ assigns a positive weight for each edge in E, and a constant $R > 0$, a vertex $u \in V$ is said to* dominate *another vertex $v \in V$ if $d(u, v) + d(v, u) \leq 2R$, where d denotes the directed path weight. The* GDSP *problem computes a set $D \subseteq V$ of* minimal *cardinality such that for any $v \in V - D$, there exists a vertex $u \in D$ such that u* dominates *v.*

GDSP is NP-hard due to a direct reduction from DSP where all edge weights are assumed to be 1, and $R = 2$.

4.1.2 Greedy Algorithm for GDSP

To solve GDSP, we adapt the greedy algorithm proposed in [10]. We refer to our algorithm as *Greedy-GDSP*. The only input parameter to the clustering process is the cluster radius R.

First, the *dominating sets* of every vertex v, denoted by $\Lambda(v)$, are computed. This is achieved by running the shortest path algorithm from a source vertex till distance $2R$. The dominance relationship is *symmetric*, i.e., $u \in \Lambda(v) \Leftrightarrow v \in \Lambda(u)$.

Symbol	Description
$TC(s)$	Set of trajectories covered by site s
$\mathcal{CL}(g)$	Neighbors of cluster g
$\mathcal{TL}(g)$	Set of trajectories passing through cluster g
$TC(g)$	Set of trajectories passing through g and its neighbors $\mathcal{CL}(g)$
$\widehat{d_r}(T, s)$	Estimate of $d_r(T, s)$ in the clustered space
$\widehat{TC}(s)$	Estimate of $TC(s)$ in the clustered space
γ	Resolution at which index instances change
t	Number of index instances
$\mathcal{I}_p$	Index instance
R_p	Cluster radius for $\mathcal{I}_p$
η_p	Number of clusters for $\mathcal{I}_p$
$\Lambda(v)$	Dominating set for node v
f	Number of bit vectors for FM sketch
ϵ	Error parameter for FM sketch

Table 4: Important notations used in the algorithms.

The main part of the Greedy-GDSP algorithm is iterative. In the first iteration, the vertex v that dominates the largest number of vertices is chosen. The set of dominated vertices, $\Lambda(v)$, forms a new *cluster* with v as the cluster center. The vertices v and $\Lambda(v)$ are not considered for further comparisons. In addition, the vertices in $\Lambda(v)$ are removed from the dominating sets of other non-clustered vertices. In other words, for each $u \in \Lambda(v)$, if $u \in \Lambda(w)$ for some non-clustered vertex w then $\Lambda(w) = \Lambda(w) - \{u\}$. In the subsequent iterations, the vertex with the largest *incremental* dominating set as produced from previous iterations is chosen. The dominated vertices that are not part of other clusters form a new cluster. The algorithm terminates when all the vertices are clustered.

Similar to the INC-GREEDY algorithm, we use FM sketches to efficiently update the dominating sets and choose the vertex with the largest incremental dominating set in each iteration. The details are same as those described in Sec. 3.5 with trajectory covering sets for each candidate site in INC-GREEDY replaced by dominating sets for each vertex in Greedy-GDSP.

4.1.3 Analysis of Greedy-GDSP

The next two theorems characterize the approximation bound and time complexity of Greedy-GDSP.

THEOREM 5. *The cardinality of the dominating set computed using the proposed algorithm is within an approximation bound of $(1 + \epsilon').(1 + \ln n)$ of the optimal where ϵ' is the approximation error of FM sketches.*

PROOF. Following the analysis in [10], the Greedy-GDSP algorithm offers an approximation bound of $H_n = 1 + \frac{1}{2} + \cdots + \frac{1}{n} \leq 1 + \ln n$ which was shown to be tight unless $P = NP$. This, however, does not consider the approximation due to the use of FM sketches. Incorporating that, the bound becomes $(1+\epsilon').(1+\ln n)$ where ϵ' is the approximation error of FM sketches. □

THEOREM 6. *Greedy-GDSP runs in $O(|V|.(\nu \log \nu + \eta))$ time where ν is the maximum number of vertices that are reachable within the largest round-trip distance R_{max} from any vertex v, and η is the number of clusters returned by the algorithm.*

The proof is given in Appendix B.2.

4.2 Selection of Cluster Representatives

In order to run INC-GREEDY on the clusters, each cluster needs to choose a *representative* candidate site. This may be different from the cluster center that was used to construct the cluster. The flexibility is needed since the cluster representative should *necessarily* be a candidate site, although the cluster center may be any vertex in V. Taking into account the fact that the cluster representative should summarize the information about the cluster and the trajectories that pass through it, and use this information to compete against the other cluster representatives in the online phase, we study two alternatives of choosing the cluster representative:

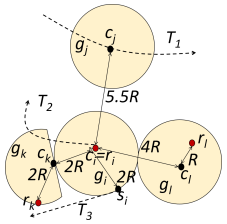

(a) Example

Cluster ID	CL	TL	TC
g_i	g_j, g_k, g_l	T_2, T_3	T_1, T_2, T_3
g_j	g_i	T_1	T_1, T_2, T_3
g_k	g_i	T_2	T_2, T_3
g_l	g_i	-	T_2, T_3

(b) Covering Sets

Figure 3: NETCLUS example using $\gamma = 0.5$.

1. The *most frequently* accessed candidate site, i.e., the one through which the largest number of trajectories pass through.

2. The candidate site that is *closest* to the cluster center.

While the first option guarantees that the utility of the cluster is at least that of its best site, the second summarizes the distribution of trajectories better. Empirical studies show that the utilities returned by the two alternatives are quite similar, but the second alternative is marginally better. Consequently, we adopt the second option.

4.3 Cluster Information

Suppose the above clustering algorithm produces clusters of radius R, i.e., the maximum round-trip distance of any node within a cluster to its cluster center is at most $2R$. Then, a pair of clusters are considered as *neighbors* of each other if their centers are within a round-trip distance of $4.R(1+\gamma)$, where $\gamma \in (0, 1]$ is the index resolution parameter to be described in Sec. 4.4. This choice of neighborhood is explained in Sec. 5.1.

As part of the index structure, every cluster g_i stores the following information:

1. Cluster *center*, c_i.

2. Cluster *representative*, r_i.

3. Trajectory set, i.e., list of *trajectories* passing through at least one site in g_i, along with their round-trip distance to c_i, $\mathcal{TL}(g_i) = \langle T_j, d_r(T_j, c_i)\rangle$.

4. Cluster *neighbors* along with the round-trip distance between their centers and c_i, $\mathcal{CL}(g_i) = \{\langle g_j, d_r(c_i, c_j)\rangle\}$, sorted by $d_r(c_i, c_j)$.

5. Set of *nodes* in the cluster and their round-trip distance to c_i, $\{\langle v_i, d_r(v_i, c_i)\rangle\}$.

The set $\mathcal{TL}(g_i)$ is computed by scanning the sequence of nodes in each trajectory T_j. If $v_i \in T_j$, then T_j is added to $\mathcal{TL}(g(v_i))$, where $g(v_i)$ denotes the cluster in which v_i resides. Thus, a trajectory is represented as a *sequence of clusters*. As neighboring nodes in any trajectory are likely to fall into the same cluster, this allows a *compressed* representation of the trajectory by collapsing consecutive copies of the same cluster into one. This compression contributes towards the efficiency of NETCLUS in terms of both space and time.

EXAMPLE 2. Fig. 3b illustrates an example of NETCLUS clustering with cluster radius R and $\gamma = 0.5$. The clusters g_i, g_j, g_k and g_l have centers c_i, c_j, c_k and c_l respectively. The cluster g_i has two candidate sites r_i and s_i. Since r_i is located at c_i, it is chosen as the cluster-representative. While cluster g_j has no candidate site, each of the clusters g_k and g_l have one candidate site, namely, r_k and r_l, respectively, each of which is a cluster-representative. The distance between the cluster centers are as follows: $d_r(c_i, c_j) = 5.5R$, $d_r(c_i, c_k) \gtrsim 2R$, and $d_r(c_i, c_l) = 4R$ where $\gtrsim$ means just greater than. The distance between any other pair of cluster centers is greater than or equal to $6R$. Given that $\gamma = 0.5$, the distance between any two neighboring cluster centers lies in the range $[2R, 6R)$. Based on this, the cluster neighbors, $\mathcal{CL}$, are shown in the figure. The trajectories T_1, T_2 and T_3 pass through nodes c_j, r_i and s_i, respectively. The figure lists the trajectory sets $\mathcal{TL}$ for each cluster. Note that when $\tau \geq 4R$, it is guaranteed that any site covers any trajectory that passes through the same cluster. For example, r_k covers T_2 as it passes through the cluster g_k that contains r_k. □

4.4 Multi-Resolution Index Structure

We next explain how the multi-resolution index structure, NETCLUS, is built by using the clustering algorithm outlined above. Assume that the normal range of query coverage threshold τ is $[\tau_{min}, \tau_{max})$. (We discuss the two extreme cases later.) NETCLUS maintains t instances of index structures $\mathcal{I}_0, \ldots, \mathcal{I}_{t-1}$ of varying cluster radii. From one instance to the next, the radius increases by a factor of $(1+\gamma)$ for some $\gamma > 0$. Thus, the total number of index instances is $t = \lfloor \log_{(1+\gamma)}(\tau_{max}/\tau_{min})\rfloor + 1$. For each instance, all the clusters and their associated information are stored.

Consider a particular index instance $\mathcal{I}_p$ with cluster radius R_p. As discussed above, the maximum round-trip distance from a site s_i belonging to the cluster g_i to a trajectory T_j that passes through g_i is at most $d_r(T_j, s_i) \leq 4.R_p$. Thus, if the coverage threshold $\tau < 4.R_p$, then it is not guaranteed if s_i covers T_j or not. Hence, the index instance $\mathcal{I}_p$ is not useful for any $\tau < 4.R_p$ and a finer instance with a lesser cluster radius should be used.

On the other hand, if τ is too large, too many neighboring clusters may cover a trajectory. Therefore, intuitively, it makes sense to switch to a higher index instance with a larger cluster radius so that less number of clusters need to be processed. The parameter $(1+\gamma)$ captures the ratio of τ to $4.R_p$ beyond which the switch is made. Thus, if $\tau > 4.R_p.(1+\gamma)$, a higher index instance is used.

Therefore, the range of useful τ for the index instance $\mathcal{I}_p$ is $[4.R_p, 4.R_p.(1+\gamma))$. Hence, the lowest cluster radius is $R_0 = (\tau_{min}/4)$, and the successive cluster radii for instances $\mathcal{I}_p$, $p = 1, \ldots, t-1$ are $R_p = (1+\gamma)^p R_0$. From one index instance to the next, as the cluster radius R_p grows, the number of clusters, $|\mathcal{I}_p|$, falls exponentially.

Choice of γ: The number of index instances t depends on γ. A smaller value of γ creates more number of instances, thereby requiring larger storage and offline running time. The approximation error is also affected by γ. When γ is smaller, the range of τ handled by a particular index instance is tighter. Therefore, the distance approximations are better. Experimental results showing the empirical impact of γ are discussed in Sec. 8.2.

Extreme cases of τ: The extreme values of the range of τ, namely, τ_{min} and τ_{max}, are assigned respectively as the *minimum* and *maximum* round-trip distance between any two sites in $\mathcal{S}$. This particular choice is guided by the following analysis. If there is a query with $\tau < \tau_{min}$, then the method degenerates to normal INC-GREEDY as each site becomes a cluster by itself. If, on the other hand, $\tau \geq \tau_{max}$, then NETCLUS reports any k sites, as each site covers every other site, and consequently, all the trajectories. Hence, the multi-resolution NETCLUS is applicable to *all* query coverage thresholds.

5. QUERYING USING NETCLUS

We next explain the online phase of querying that starts after the query parameters (k, τ) are available.

The first important consideration is choosing the index instance $\mathcal{I}_p$ that supports the given query threshold τ. The index p is computed as $p = \lfloor \log_{(1+\gamma)}(\tau/\tau_{min})\rfloor$. This ensures that $4.R_p \leq \tau < 4.R_p.(1+\gamma)$ where R_p is the cluster radius for $\mathcal{I}_p$.

We next discuss how to apply TOPS on the clustered space.

$\widehat{d}_r$	r_i	r_k	r_l
T_1	$5.5R$	$\gtrsim 9.5R$	$10.5R$
T_2	$0R$	$\leq 4R$	$5R$
T_3	$2R$	$\gtrsim 6R$	$7R$

Table 5: Distance estimates $\widehat{d}_r(\cdot)$ for example in Fig. 3a.

5.1 TOPS-Cluster Problem

Consider a cluster g_i with its representative r_i, and a trajectory T_j passing through a cluster g_j, where g_j may or may not be equal to g_i. Then $T_j \in TC(r_i)$ if and only if $d_r(T_j, r_i) \leq \tau$. In the clustered space, however, we only store the distances of each trajectory from the centers of the clusters that it passes through. Hence, it is not possible to compute $d_r(T_j, r_i)$ without extensive online computation. Hence, an *approximate* distance $\widehat{d}_r(T_j, r_i)$ is computed and used. The round-trip distance estimate from T_j to r_i is

$$\widehat{d}_r(T_j, r_i) = d_r(T_j, c_j) + d_r(c_j, c_i) + d_r(c_i, r_i) \tag{9}$$

It is important to note that the distance can be estimated using *only* the information computed in the offline phase. Since the distances are approximate, the *approximate trajectory cover* of r_i is

$$\widehat{TC}(r_i) = \{T_j \in TC(g_i) | \widehat{d}_r(T_j, r_i) \leq \tau\} \tag{10}$$

where $TC(g_i) = \mathcal{TL}(g_i) \cup \{T_j \in \mathcal{TL}(g_j) | g_j \in \mathcal{CL}(g_i)\}$ consists of the trajectories passing through g_i and its neighbors $\mathcal{CL}(g_i)$.

For any $T_j \in \widehat{TC}(r_i)$, $T_j \in TC(g_i)$, if and only if there exists a cluster g_j such that $T_j \in \mathcal{TL}(g_j)$ and $d_r(c_i, c_j) \leq \tau$. This follows from the fact that $\widehat{d}_r(T_j, r_i) \leq \tau$. For the index instance $\mathcal{I}_p$, since $\tau \leq 4R_p(1+\gamma)$, therefore, this condition reduces to $d_r(c_i, c_j) \leq 4R_p(1+\gamma)$. This is the reason why the neighborhood of a cluster is defined as those whose centers are within a round-trip distance of $4R_p(1+\gamma)$ in Sec. 4.3.

Consequently, to compute the set $\widehat{TC}(r_i)$, it is sufficient to examine only the trajectory sets of the neighbors of the cluster g_i. For each trajectory $T_j \in \mathcal{TL}(g_j)$ where g_j is a neighbor of g_i, the approximate distance $\widehat{d}_r(T_j, r_i)$ is computed. The trajectory T_j is included in $\widehat{TC}(r_i)$ if $\widehat{d}_r(T_j, r_i) \leq \tau$.

EXAMPLE 3. The TC sets for the three clusters in Fig. 3a are shown in Fig. 3b. Using Eq. (9), the distance estimates between each pair of cluster representative and trajectory is shown in Table 5. Since $\gamma = 0.5$, therefore, the supported range of τ is $[4R, 6R)$. Thus, if $\tau = 4R$, $\widehat{TC}(r_i) = \{T_2, T_3\}$, $\widehat{TC}(r_k) = \{T_2\}$, and $\widehat{TC}(r_l) = \varnothing$. Similarly, if $\tau = 5.75R$, $\widehat{TC}(r_i) = \{T_1, T_2, T_3\}$, $\widehat{TC}(r_k) = \{T_2\}$, and $\widehat{TC}(r_l) = \{T_2\}$. □

If a trajectory $T_j \in \widehat{TC}(r_i)$, then it also lies in the set $TC(r_i)$ since $d_r(T_j, r_i) \leq \widehat{d}_r(T_j, r_i) \leq \tau$. However, the reverse is not true, since there may be a trajectory T_j such that $d_r(T_j, r_i) \leq \tau$, but the estimate $\widehat{d}_r(T_j, r_i) > \tau$. Therefore, $\widehat{TC}(r_i) \subseteq TC(r_i)$. For example, in Fig. 3a, $d_r(T_3, r_k) \leq 4R$, but $\widehat{d}_r(T_3, r_k) \gtrsim 6R$ which exceeds any supported value of τ. Thus, $T_3 \notin \widehat{TC}(r_k)$.

Finally, based on the query preference function ψ, the preference score $\psi(T_j, r_i)$ is computed between all cluster representatives r_i and their trajectory covers $T_j \in \widehat{TC}(r_i)$.

Using these approximate covering sets, we run the following instance of TOPS problem, called TOPS-CLUSTER.

PROBLEM 3 (TOPS-CLUSTER). *Given an index instance $\mathcal{I}_p$ defined over the road network $G = (V, E)$, suppose $\widehat{\mathcal{S}} \subseteq S$ denote the set of cluster representatives in $\mathcal{I}_p$. TOPS-Cluster problem seeks to report a set of k cluster representatives, $\mathcal{Q} \subseteq \widehat{\mathcal{S}}$, $|\mathcal{Q}| = k$, such that $U(\mathcal{Q})$ is* maximal.

To solve TOPS-Cluster, we employ INC-GREEDY on the set of cluster representatives $\widehat{\mathcal{S}}$ using the above covering sets $\widehat{TC}(r_i)$.

When the preference function ψ is binary, FM sketches can be employed for faster updating of marginal utilities during the execution of INC-GREEDY on the cluster representatives, in the same manner as described in Sec. 3.5.

5.2 Analysis of NetClus

Quality Analysis: The first result is due to a direct application of Lem. 2.

COROLLARY 1. *The utility of the set $\widehat{\mathcal{Q}}$ returned by the* NETCLUS *framework is bounded as follows: $U(\widehat{\mathcal{Q}}) \geq (k/|\widehat{\mathcal{S}}|)U(\widehat{\mathcal{S}})$.*

If the index instance $\mathcal{I}_p$ is used for a particular query threshold τ, then $|\widehat{\mathcal{S}}|$ is at most the number of clusters, η_p. The next result states the approximation guarantees offered by the NETCLUS framework.

THEOREM 7. *The approximation bound offered by* NETCLUS *for the binary instance of TOPS is (k/η_p). For a general preference function $\psi(T_j, s_i) = f(\cdot)$, the approx. bound is $f(\tau).(k/\eta_p)$.*

PROOF. Assuming all nodes are candidate sites, we observe that each trajectory $T_j \in \mathcal{T}$ is covered by the cluster representative r_i of a cluster g_i that it passes through. This is because the maximum round-trip distance between T_j and r_i is at most $4R_p$ and τ is at least $4R_p$.

This ensures that $U(\widehat{\mathcal{S}}) = \sum_{T_j \in \mathcal{T}} U_j = |\mathcal{T}| = m$ because $\forall T_j \in \mathcal{T},\ U_j = 1$.

Since the maximum utility of the optimal algorithm for TOPS can be at most m, following the result in Cor. 1, the approximation bound is at least k/η_p where $\eta_p = |\widehat{\mathcal{S}}|$.

Next, consider a general preference function $\psi(T_j, s_i) = f(\cdot)$ where f is a positive non-increasing function of $d_r(T_j, s_i)$ such that $f(0) = 1$. If T_j passes through the cluster g_i, then $d_r(T_j, r_i) \leq \tau$. Hence, $U_j = \max\{\psi(T_j, r_i) | r_i \in \widehat{\mathcal{S}}\} \geq f(\tau)$ since f is non-increasing. Therefore, $U(\widehat{\mathcal{S}}) = \sum_{T_j \in \mathcal{T}} U_j \geq f(\tau).m$. Since the preference scores lie in the range $[0, 1]$, the utility offered by any optimal algorithm for TOPS is at most m. Therefore, following Cor. 1, the approximation bound is at least $f(\tau).(k/\eta_p)$. □

To solve the binary instance of TOPS problem, the FM sketches may be employed while running the INC-GREEDY algorithm on the cluster representatives. The resulting scheme is referred to as FM-NETCLUS. In that case, the bound is updated as follows.

THEOREM 8. *The approximation bound of* FM-NETCLUS *for the binary instance of TOPS is $(k/\eta_p).(1+\epsilon)^k$, where ϵ is the error parameter provided by the FM sketch.*

PROOF. If the error parameter for FM sketch is ϵ, running it for k iterations produces an error bound of at most $(1+\epsilon)^k$. In conjunction with Th. 7, the required error bound is obtained. □

Complexity Analysis: For a given value of τ, suppose the index instance $\mathcal{I}_p$ with $|\mathcal{I}_p| = \eta_p$ clusters is used. Assume that the largest number of trajectories passing through a cluster is $\xi_p = \max\{|\mathcal{TL}(g_i)|\}$, and λ_p is the largest number of vertices in any cluster in $\mathcal{I}_p$.

THEOREM 9. *The time and space complexities of* NETCLUS *are $O(k.\eta_p.\xi_p)$ and $O(\sum_{p=1}^{t}(\eta_p(\xi_p + \lambda_p)))$ respectively.*

The proof is provided in Appendix B.3.

The average values of η_p, λ_p and $|\mathcal{TL}|$ are shown in Table 11, for different values of cluster radius R_p.

6. HANDLING DYNAMIC UPDATES

In this section, we discuss how the NETCLUS framework efficiently handles dynamic updates of trajectories and candidate sites. We assume that the underlying road network does not change. In each of the following cases, the updates are processed for *all* the index instances (of varying cluster radii).

Addition of a site: Suppose a location s_{add} is identified as a new candidate site, i.e., s_{add} gets added to $\mathcal{S}$. If s_{add} is already in V, its cluster g_{add} is identified. Otherwise, it is added to the cluster g_{add} whose cluster center c_{add} is the closest. To determine the closest cluster center, the neighbors $N(s_{add})$ of s_{add} in G are used. The round-trip distance to a cluster center c_i is estimated using $\min_{s_l \in N(s_{add})}\{d_r(s_{add}, s_l) + d_r(s_l, c_i)\}$ if $d_r(s_l, c_i)$ is available. Suppose c_{near} is the nearest cluster center to s_{add}. If the distance $d_r(s_{ad}, g_{nearest}) > 2R_p$, then we create a new cluster g_{add} with s_{add} as its center. If the identified cluster g_{add} does not have a cluster-representative, then s_{add} is marked as its new cluster representative. Else, it is determined if s_{add} can be a better representative for g_{add} as discussed in Sec. 4.2. Finally, the exact round-trip distance to the cluster center, $d_r(s_{add}, c_{add})$, is computed.

Deletion of a site: Suppose a particular site s_{del} is no longer viable for a service and, therefore, needs to be deleted from S. Suppose, s_{del} lies in the cluster g_{del}. First, it is untagged as a candidate site in g_{del}. If s_{del} is not the cluster representative of g_{del}, nothing more needs to be done. Otherwise, another candidate site, if available, is chosen as the new cluster representative using the methodology described in Sec. 4.2.

Addition of a trajectory: Suppose a new trajectory T_{add} is added. It is first mapped into a sequence of clusters, $g_1, \ldots, g_l$. For each such cluster g_i, T_{add} is added to the set $\mathcal{TL}(g_i)$ and its round-trip distance to the cluster center c_i of g_i, $d_r(T, c_i)$, is computed and stored. In addition, g_i is added to the set $CC(T_{add})$. The procedure is essentially the same one discussed in Sec. 5.

Deletion of a trajectory: Suppose a trajectory T_{del} is deleted. Assume that the coverage set of T_{del} is $CC(T_{del}) = \{g_1, \ldots, g_l\}$. For each such cluster g_i, T_{del} is removed from its coverage set $\mathcal{TL}(g_i)$. Finally, the set $CC(T_{del})$ is deleted.

While multiple updates can be applied one after another, batch processing is more efficient. Sec. 8.8 shows that the updates are handled quite efficiently.

7. EXTENSIONS AND VARIANTS OF TOPS

In this section, we present a few extensions and variants within the TOPS framework and discuss how the INC-GREEDY algorithm for TOPS can be adapted to solve these problems. As NETCLUS essentially runs INC-GREEDY on the cluster representatives, it can also be adapted in a similar manner.

7.1 Cost Constrained TOPS

In this problem, referred to as TOPS-COST, each site $s_i \in \mathcal{S}$ has a cost $cost(s_i)$ associated with it, and the goal is to select a set of sites within a fixed budget B such that the sum of trajectory utilities is maximized. Formally, the problem is stated as follows.

PROBLEM 4 (TOPS-COST). *Given a set of trajectories $\mathcal{T}$, a set of candidate sites $\mathcal{S}$ where each site $s_i \in \mathcal{S}$ has a fixed cost $cost(s_i)$,* TOPS-COST *problem with query parameters (B, τ, ψ) seeks to report a set $\mathcal{Q} \subseteq \mathcal{S}$, that maximize the sum of trajectory utilities, i.e.,* $\max_{\mathcal{Q} \subseteq \mathcal{S}} U(\mathcal{Q}) = \sum_{T_j \in \mathcal{T}} U_j$ *such that* $U_j = \max_{s_i \in \mathcal{Q}}\{\psi(T_j, s_i)\}$, $cost(\mathcal{Q}) = \sum_{s_i \in \mathcal{Q}} cost(s_i) \leq B$.

The TOPS problem reduces to TOPS-COST by assigning unit cost to each site and $B = k$. In contrast to TOPS, TOPS-COST does not restrict the number of sites selected in the answer set. Since TOPS is NP-hard, so is TOPS-COST.

The INC-GREEDY algorithm can be adapted based on the greedy heuristic for the budgeted maximum coverage problem [24] to solve TOPS-COST. The algorithm starts with an empty set of sites $\mathcal{Q} = \varnothing$ and proceeds in iterations. In each iteration, it selects a site $s_i \in \mathcal{S} - \mathcal{Q}$ such that $(U(\mathcal{Q} \cup \{s_i\}) - U(\mathcal{Q}))\,/cost(s_i)$ is *maximal*. If $cost(s_i)$ is within the remaining budget $B - cost(\mathcal{Q})$, it is added to $\mathcal{Q}$; otherwise, it is pruned from $\mathcal{S}$. This process continues until $\mathcal{S} = \varnothing$.

It was shown in [24] that the above approach can perform arbitrarily bad. Thus, in order to bound the approximation guarantee, the algorithm is augmented with the following step. Assume s_{max} to be a candidate site such that $cost(s_{max}) \leq B$ and $U(\{s_{max}\})$ is maximal. The algorithm returns either the site s_{max} or the set $\mathcal{Q}$ whichever offers the maximum utility. Following the analysis in [24], this scheme is guaranteed to produce a solution with an approximation bound of $(1 - 1/e)/2$.

7.2 Capacity Constrained TOPS

In this problem, referred to as TOPS-CAPACITY, each site $s_i \in \mathcal{S}$ has a fixed capacity $cap(s_i)$ that denotes the maximum number of trajectories it can serve. The goal is to select a set $\mathcal{Q} \subseteq \mathcal{S}$ of size k such that the sum of trajectory utilities is maximized. Formally, the problem is stated as follows.

PROBLEM 5 (TOPS-CAPACITY). *Consider a set of trajectories $\mathcal{T}$, a set of candidate sites $\mathcal{S}$ where each site $s_i \in \mathcal{S}$ can serve at most $cap(s_i)$ trajectories. For any set $\mathcal{Q} \subseteq \mathcal{S}$, let x_{ji} be a boolean indicator variable such that $x_{ji} = 1$ if and only if the trajectory T_j can be served by the site $s_i \in \mathcal{Q}$.* TOPS-CAPACITY *problem with query parameters (k, τ, ψ) seeks to report a set $\mathcal{Q} \subseteq \mathcal{S}, |\mathcal{Q}| = k$, that maximizes the sum of trajectory utilities, i.e.,* $\max_{\mathcal{Q} \subseteq \mathcal{S}} U(\mathcal{Q}) = \sum_{T_j \in \mathcal{T}} U_j$ *such that* $U_j = \max_{s_i \in \mathcal{Q}}(\psi(T_j, s_i).x_{ji})$, *and* $\forall s_i \in \mathcal{Q}$, $\sum_{T_j \in \mathcal{T}} x_{ji} \leq cap(s_i)$.

TOPS reduces to TOPS-CAPACITY by assigning the capacity of each site to be infinite or more than the total number of trajectories in the dataset. Hence, TOPS-CAPACITY is also NP-hard.

INC-GREEDY can be adapted to solve TOPS-CAPACITY in the following manner. The algorithm starts with an empty set $\mathcal{Q}_0 = \varnothing$ and sets $\forall T_j \in \mathcal{T}, U_j = 0$. Suppose the set of sites selected after iteration $\theta = 1, \ldots, k$, is denoted by $\mathcal{Q}_\theta$. In each iteration θ, it augments $\mathcal{Q}_{\theta-1}$ by selecting a site $s_\theta \in \mathcal{S} - \mathcal{Q}_{\theta-1}$ that offers the maximal marginal gain in utility.

It then updates the trajectory utilities U_j. The utility of T_j due to $\mathcal{Q}_\theta$ is $U_j(\mathcal{Q}_\theta) = \max_{s_i \in \mathcal{Q}_\theta} \psi(T_j, s_i)$. The marginal gain in the utility of T_j due to addition of s_θ is $U_j(s_\theta) = U_j(\mathcal{Q}_{\theta-1} \cup \{s_\theta\}) - U_j(\mathcal{Q}_{\theta-1})$. Since any site $s_i \in \mathcal{S} - \mathcal{Q}_{\theta-1}$ can serve at most $\alpha_i = \min\{|TC(s_i)|, cap(s_i)\}$ trajectories, its marginal utility is defined as the sum of the largest α_i trajectory marginal utilities $U_j(s_i)$.

Since the objective function of the TOPS-CAPACITY problem is identical to that of TOPS, it follows the non-decreasing submodular property. Thus, INC-GREEDY offers the same approximation bound of $1 - 1/e$, as stated in Lem. 1.

7.3 TOPS with Existing Services

Optimal location queries usually factor in existing service locations before identifying new ones. The problem is NP-hard. The INC-GREEDY algorithm can take into account the existing service locations as follows.

Suppose $\mathcal{E}_\mathcal{S}$ is the set of existing service locations. On receiving the query parameter τ, the covering sets TC and SC over the set of sites $\mathcal{S} \cup \mathcal{E}_\mathcal{S}$ are computed. INC-GREEDY starts with $\mathcal{Q}_0 = \mathcal{E}_\mathcal{S}$ and updates the marginal utilities of the sites in $\mathcal{S}$. The remaining algorithm stays unaltered. The algorithm terminates after selecting k sites from the set $\mathcal{S}$, in the same manner as TOPS.

An advantage and important feature of INC-GREEDY is that the site chosen in a given iteration depends solely on what the existing service locations are, and not on how they were chosen.

Since the initial utility $U(\mathcal{Q}_0) = U(\mathcal{E}_\mathcal{S}) \neq 0$, the approximation bound of $1-1/e$ is not directly applicable. However, we next show that the same bound holds.

For any set $\mathcal{Q} \subseteq \mathcal{S}$, the extra utility, defined as $U'(\mathcal{Q}) = U(\mathcal{Q}) - U(\mathcal{E}_\mathcal{S})$ where $U(\mathcal{E}_\mathcal{S})$ is the utility offered by the existing services, is non-negative. Now we have $U'(\mathcal{Q})$ as a non-decreasing submodular function with $U'(\mathcal{Q}_0) = 0$. Let OPT and $\mathcal{Q}$ denote the set of sites returned by an optimal algorithm and INC-GREEDY respectively. Then, following Lem. 1, $U'(\mathcal{Q}) \geq (1-1/e)U'(\text{OPT})$. This leads to $U(\mathcal{Q}) \geq (1-1/e)U(\text{OPT}) + U(\mathcal{E}_\mathcal{S})/e$. Since $U(\mathcal{E}_\mathcal{S}) \geq 0$, hence, $U(\mathcal{Q}) \geq (1-1/e)U(\text{OPT})$.

7.4 Other TOPS Variants

There are certain variants of trajectory-aware service location problems that already exist in the literature. The proposed TOPS formulation generalizes these variants and, thus, the INC-GREEDY algorithm can be used to solve them. We next discuss some of the important variations.

TOPS1: Binary world: This is the simple binary instance of TOPS query that is defined in Def. 3. The problem is still NP-hard [2]. INC-GREEDY offers the same approximation bound for TOPS1, as in the case of TOPS (Th. 3).

TOPS2: Maximize market size: Instead of operating in a binary world where a trajectory is either covered or not covered, in this formulation, the aim is to maximize the probability of a trajectory being covered [2]. The probability $p(T_j, s_i)$ of T_j being covered by site s_i is modeled as a convex function of the distance $d_r(T_j, s_i)$ between them. This problem aims to set up k services that maximizes the expected number of total trajectories served. This is a special case of our proposed TOPS formulation with $\psi(T_j, s_i) = p(T_j, s_i)$ if $d_r(T_j, s_i) \leq \tau$ and 0 otherwise. It is again NP-hard [2]. INC-GREEDY offers the same approximation bound for TOPS2, as in the case of TOPS (Th. 3).

TOPS3: Minimize user inconvenience: Assuming that each user on its trajectory would necessarily avail a service, this problem aims to minimize the expected deviation incurred by a user [2]. This can be handled through the TOPS framework by setting the preference score to $\psi(T_j, s_i) = -d_r(T_j, s_i)$, and τ to ∞. Since the trajectory utility is defined to be the maximum of the preference scores of the selected sites, maximizing the sum of trajectory utilities would minimize the total user deviation. TOPS3 is NP-hard owing to reduction from the k-medians problem [2]. The approximation bound of INC-GREEDY for this problem is not yet known.

TOPS4: Best service locations under fixed market share: The aim here is to place the minimum number of services that can capture a fixed share of the market comprising of the user trajectories [7]. The problem is complementary to TOPS1 and asks the following: What is the smallest set of sites $\mathcal{Q} \subseteq \mathcal{S}$ such that at least β fraction of $|\mathcal{T}|$ is covered, where $0 < \beta \leq 1$? This problem is NP-hard [7]. Since INC-GREEDY algorithm is iterative, it selects as many sites that are necessary to cover the desired fraction of users. Note that the set cover problem directly reduces to this problem, and so does the greedy heuristic for the set cover problem to INC-GREEDY. Hence, INC-GREEDY algorithm offers the same approximation bound of $1 + \ln n$ for TOPS4.

7.5 Generic Framework

From the above discussions, it is clear that the TOPS framework is highly generic and can absorb many extensions and variations with little or no modification in the prposed algorithms INC-GREEDY and NETCLUS. Also, importantly, the framework also enables *combining* multiple extensions and variants. For example, TOPS-COST and TOPS-CAPACITY extensions can be merged to create a new version of TOPS. In Sec. 8.7, we discuss experimental evaluation of some of the above extensions and variants.

Dataset	Type	#Trajectories	#Sites
Beijing-Small	Real	1,000	50
Beijing	Real	123,179	269,686
Bangalore	Synthetic	9,950	61,563
New York	Synthetic	9,950	355,930
Atlanta	Synthetic	9,950	389,680

Table 6: Summary of datasets.

8. EXPERIMENTAL EVALUATION

In this section, we perform extensive experiments to establish (1) *Efficiency:* that NETCLUS is efficient, practical and scales well to real-life datasets, and (2) *Quality:* that NETCLUS produces solutions that are close to those of INC-GREEDY, which serves as the *baseline* technique.

Most of the results are shown for the binary version of the problem since it is easier to comprehend. Sec. 8.7 shows the results for various extensions and variants of TOPS.

The experiments were conducted using Java (version 1.7.0) platform on an Intel(R) Core i7-4770 CPU @3.40GHz machine with 32 GB RAM running Ubuntu 14.04.2 LTS OS.

8.1 Evaluation Methodology

Algorithms: We evaluate the performance of three different algorithms to address TOPS: OPT, INC-GREEDY, and NETCLUS, which are henceforth referred to as OPT, INCG, and NETCLUS respectively in text and in the figures. The variants of INCG and NETCLUS, based on the FM sketches, are henceforth referred to as FMG and FMNETCLUS respectively.

To the best of our knowledge, there is no existing algorithm for TOPS that works on real trajectories on city-scale road networks. The state-of-the-art algorithm for TOPS1 problem, proposed in [2], when adapted for our case, can be reduced to INC-GREEDY. Hence, INC-GREEDY acts as the baseline algorithm for evaluation on TOPS1.

Variants: The specific variant of TOPS on which we evaluated most of the experiments was the binary instance defined in Def. 3 or TOPS1. We choose to evaluate this particular variant of TOPS due to three reasons: (1) Usually, such binary optimization problems are the worst case instances of the general integer optimization problems and are, therefore, hardest to approximate. (2) The site preference function is simple. (3) It has several applications in transportation science and operations research. However, we also evaluate a few other TOPS variants and extensions in Sec. 8.7.

Metrics of evaluation: The main metrics of evaluation were (a) total *utility* measured as a percentage of the total number of trajectories m, and (b) query *running time*. The two basic parameters studied were (i) number of service locations k, and (ii) coverage threshold τ, whose default values were 5 and 0.8 Km. respectively.

We conducted experiments on both real and synthetic datasets, whose details are shown in Table 6. For simplicity, we assume that the number of candidate sites is the same as the number of nodes in the graph, unless otherwise stated.

Real datasets: We used GPS traces of taxis from Beijing consisting of user trajectories generated by tracking taxis for a week [41,42]. This is the most widely used and one of the largest publicly available trajectory datasets. To generate trajectories as sequences of road intersections, the raw GPS-traces were map-matched [33] to the Beijing road network extracted from OpenStreetMap (http://www.openstreetmap.org/). The road network contains 269,686 nodes and 293,142 edges, with an underlying area of ∼1720 sq. Km.

Since TOPS is NP-hard, the optimal algorithm requires exponential time and, therefore, can be run only on a very small dataset.

γ	Time (s)	Space (GB)	Rel. Error % w.r.t. INCG
0.25	108427	14.095	3.54
0.50	3216	4.215	3.97
0.75	1652	2.374	4.53
1.00	520	1.053	5.21

Table 7: Variation across resolution of index instances, γ.

f	Utility		Rel. Error %	Time (ms)		Speed-up
	NETCLUS	FM		NETCLUS	FM	
1	47.23	26.60	43.67	846.65	16.32	51.88
2	47.23	34.27	27.45	846.65	21.66	39.09
4	47.23	39.33	16.73	846.65	32.65	25.93
10	47.23	41.27	12.62	846.65	65.32	12.96
20	47.23	43.29	8.34	846.65	116.32	7.28
30	47.23	44.96	4.81	846.65	161.53	5.24
40	47.23	45.52	3.63	846.65	216.62	3.91
50	47.23	45.89	2.84	846.65	272.18	3.11
100	47.23	46.43	1.69	846.65	984.17	0.86

Table 8: Variation across the number of FM sketches, f.

Hence, we evaluate all the algorithms against the optimal on Beijing-Small which is generated by randomly sampling 1000 trajectories from a fixed area, and then randomly selecting the set $\mathcal{S}$, consisting of 50 candidate sites, from the same area. The sampling was conducted 10 times to increase the robustness of the results. All the other experiments were done on the full Beijing dataset.

Synthetic datasets: To study the impact of city geographies, we generated three synthetic datasets that emulate trajectories in the patterns followed in New York, Atlanta and Bangalore. We used an online traffic generator tool, MNTG (http://mntg.cs.umn.edu/tg/index.php) to generate the traffic traces, that were later map-matched to generate the trajectories in the desired format.

8.2 Choice of Parameters

We first run experiments to determine the choice of two important parameters: (a) the resolution of the index instances, γ, and (b) the number of FM bit vectors, f.

As discussed earlier in Sec. 4.4, the choice of γ affects both the storage and offline run-time costs as well as quality. Table 7 lists the values for the Beijing dataset when γ is changed. The error is measured as relative loss in utility of NETCLUS with that of INCG. When γ is too small, there is almost no compression of the trajectories. As a result, the index structure size is large. On the other hand, with a very large γ, the error may be unacceptable. We fix $\gamma = 0.75$ for our experiments since it offers a nice balance of a medium sized index structure that can fit in most modern systems with an error of within 5%.

Table 8 shows how the utility and running time varies when f number of FM sketches are used as compared to the original NETCLUS. The error is measured as a relative loss in utility of FM-NETCLUS with that of NETCLUS. The values of τ and k were their default values (0.8 Km. and 5 respectively). When f is very small, the error is too large. As f increases, expectedly the error decreases while the speed-up decreases as well. When f is extremely high, the number of operations may overshadow the gains and using FM sketches may be actually slower. We fixed $f = 30$ since it produced less than 5% error with a speed-up factor of more than 5.

8.3 Comparison with Optimal

Since this integer linear program based optimal algorithm requires impractical running times, we ran it only on the Beijing-Small dataset mainly to assess the quality of the other algorithms. Fig. 4 shows that the average utility of all the algorithms are quite close to OPT although the running times are much better. (Note that the utilities in these and all subsequent figures are plotted as a percentage of the total number of trajectories.) OPT requires hours to complete even for this small dataset and, therefore, is not practical at all. Consequently, we did not experiment with OPT any further.

8.4 Quality Results

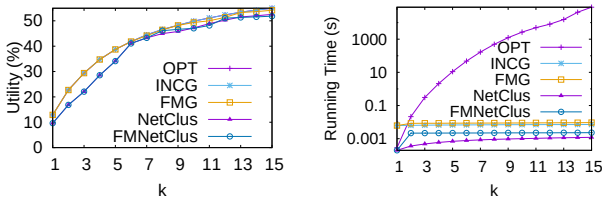

(a) Utility. (b) Running time.

Figure 4: Comparison with optimal at $\tau = 0.8$ Km.

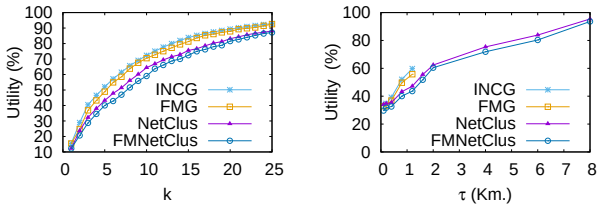

(a) Varying k (b) Varying τ

Figure 5: Quality results.

Fig. 5 shows the utility yields for different values of k and τ. The utilities of NETCLUS are close to that of INCG and are within 93% of it on an average. Owing to high memory requirements (for reasons to be discussed in the next section), INCG and FMG could not run beyond $\tau = 1.2$ Km. The utilities of FMG and FMNETCLUS are very close to that of INCG and NETCLUS, respectively.

8.5 Memory Footprint

Table 9 shows that the memory footprints of NETCLUS and FM-NETCLUS are significantly less than those of INCG and FMG. As the coverage threshold τ increases, the size of the covering sets, TC and SC, used in INCG and FMG, increase sharply. Consequently, these algorithms could not scale beyond $\tau = 1.2$ Km. On the other hand, with higher τ, NETCLUS and FMNETCLUS use lower resolution clustering instances leading to higher data compression, thereby resulting in lower memory footprints. The FM sketch based schemes require slightly more memory than their counterparts, due to storage of multiple bit vectors for each site or cluster, as applicable.

8.6 Performance Results

We next measure the performance of the algorithms for different values of k and τ. Fig. 6 shows that for $\tau \leq 1.2$ km., NETCLUS and FMNETCLUS are up to 36 times faster than INCG and FMG, respectively. For $\tau > 1.2$ km., as stated in the previous section, INCG and FMG fail to run due to high memory overheads. When τ increases, NETCLUS (and FMNETCLUS) uses a higher index instance having lesser number of clusters, leading to its efficiency. On the other hand, INCG and FMG use TC and SC covering sets of increasingly larger size, resulting in poor performance.

Note that the plots in Fig. 6a appear to be linear w.r.t. k. This is because (i) the initial cost of computing the covering sets significantly dominates the iterative phase of the algorithms, (ii) the running times are plotted in log-scale.

Although FMNETCLUS (FMG) offers a speed up of about 5 times in the algorithm running time when compared to NETCLUS (INCG respectively), its effect is negated by a relatively large initial pre-processing time required for computing the covering sets. Due to this fact, we only compare the results of NETCLUS with that of INCG in the subsequent sections.

8.7 Extensions and Variants of TOPS

We next show results of NETCLUS on different TOPS extensions and variants (discussed in Sec. 7) over the Beijing dataset.

τ (in Km.)	INCG	FMG	NETCLUS	FMNETCLUS
0.1	7.04	7.90	6.43	7.09
0.2	9.14	10.00	4.17	4.81
0.4	13.47	14.34	3.67	3.94
0.8	19.58	20.44	3.22	3.64
1.2	23.98	24.85	3.52	3.88
1.6	Out of memory		3.41	3.98

Table 9: Memory footprint of different algorithms (in GB).

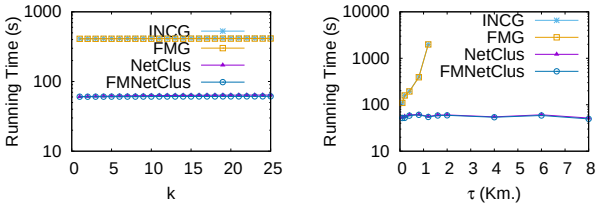

(a) Varying k (b) Varying τ

Figure 6: Running time performance.

TOPS-COST: We consider a budget of $B = 5.0$ and $\tau = 0.8$ Km. The cost of each site was assigned using a normal distribution with mean $\mu = 1.0$ and standard deviation varied between $\sigma \in [0, 1]$ (the least cost of a site was constrained to be 0.1). Fig. 7a shows that the utility increases with standard deviation (note that $\sigma = 0$ degenerates to basic TOPS). This is due to the fact that with higher standard deviation, more number of sites can be chosen with lower costs which ultimately leads to larger number of trajectories being covered. The increased number of iterations does not increase the running time much since it is a small overhead on the initial costs (Fig. 9).

TOPS-CAPACITY: We consider $k = 5$ and $\tau = 0.8$ Km. The sites were assigned varying capacities drawn from a normal distribution where the mean was varied in the range $[0.1\%, 100\%]$ of the total number of trajectories, and the standard deviation fixed at 10% of the mean. (note that mean capacity of 100% corresponds to basic unconstrained TOPS). Fig. 7b shows that, as expected, utility increases with mean capacity. NETCLUS has almost the same utility as that of INCG. We do not show the running time plots, as the algorithms for TOPS-CAPACITY are almost the same as those for TOPS and, hence, exhibit similar performance.

TOPS2: Finally, we study the TOPS2 variant where the preference function ψ was a convex function of the distance between the site and trajectory. The results, shown in Fig. 8, portray that NETCLUS has utility close to that of INCG, while being about an order of magnitude faster.

8.8 Updates of Sites and Trajectories

Table 10 shows that NETCLUS efficiently processes additions of trajectories and candidate sites over the index structure (Sec. 6). Adding a trajectory requires more time than that for a candidate site since a trajectory passes through multiple clusters in general and the covering sets, etc. of all those clusters need to be updated. Adding a site, on the other hand, requires simply finding the cluster it is in and updating the cluster representative, if applicable.

8.9 Robustness with Parameters

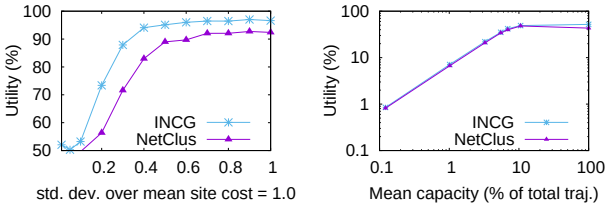

(a) Cost. (b) Capacity.

Figure 7: Utilities for TOPS extensions.

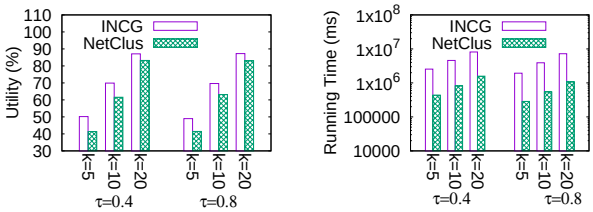

(a) Utility. (b) Running time.

Figure 8: TOPS2: A variant of TOPS.

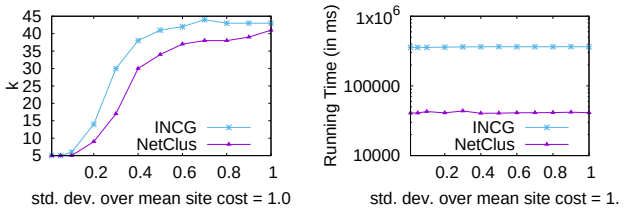

(a) Number of sites. (b) Running time.

Figure 9: TOPS under cost constraint.

Number of Sites and Trajectories: Fig. 10 shows scalability results with varying number of candidate sites and trajectories on the Beijing dataset. NETCLUS is about an order of magnitude faster than INCG.

City Geometries: We experimented with three typical city geometries, Atlanta, New York, and Bangalore (Fig. 11). New York has a star topology while Bangalore is poly-centric. Consequently, Bangalore has a larger utility percentage. Since Atlanta has a mesh structure with trajectories distributed all over the city, its utility is lowest. There is not much difference in the running times, though. Bangalore has least running time due to its smaller road network.

Length of Trajectories: To determine the effect of length of trajectories, the trajectories were divided into four classes based on their lengths and from each of the classes, 5,000 trajectories were sampled (Fig. 12). Longer trajectories are easier to cover since they pass through more number of candidate sites over a larger area and, therefore, exhibit higher utility than the shorter ones. The running time also increases with trajectory length due to more number of update operations of the marginal utilities.

8.10 Index Construction

Referring to Table 11, we observe that as the cluster radius R_p increases, the number of clusters η_p decreases as the average dominating set sizes $|\bar{\Lambda}|$ increases. Therefore, the average number of trajectories passing through a cluster $|\bar{\mathcal{TL}}|$ also increases. The average number of neighbors of a cluster $|\mathcal{CL}|$ initially increases but finally decreases. Importantly, we observe that the *offline* index construction times across different cluster radii are practical.

9. RELATED WORK

The related work falls in two main classes, *optimal location* queries [11, 14, 17, 28, 40], and *flow based facility location* problems [2–7].

An optimal location (OL) query has three inputs: a set of facilities, a set of users, and a set of candidate sites. The objective is to determine a candidate site to set up a new facility that optimizes an objective function based on the distances between the facilities and the users. Comparing OL query with TOPS query, we note that: (a) While *fixed* user locations are considered for OL queries, TOPS uses *trajectories* of users. (b) OL queries report only a single optimal location, while TOPS reports k locations. (c) Unlike OL queries that are solvable in polynomial time, TOPS is NP-hard (it is polynomially solvable only for $k = 1$).

# Trajectories added	Update time	# Candidate sites added	Update time
10000	22.83 s	10000	1.26 s
20000	44.58 s	20000	1.56 s
30000	74.07 s	30000	1.76 s
40000	92.06 s	40000	1.85 s
50000	122.69 s	50000	2.10 s

Table 10: Index update cost.

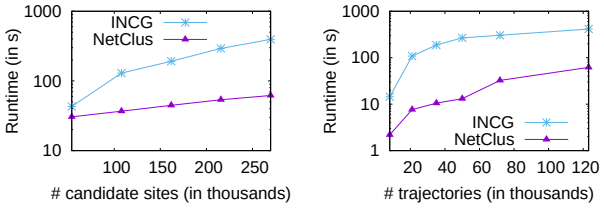

(a) Number of sites. (b) Number of trajectories.

Figure 10: Scalability results ($k = 5$ and $\tau = 0.8$ Km).

Recently, [28] studied OL queries on trajectories over a road network. Two algorithms were proposed to compute the optimal road segment to host a new service. Their work has quite a few limitations and differences when compared with our work: (a) Since a single optimal road segment is reported, their problem is polynomially solvable. (b) Their work identifies the optimal road segment, rather than the optimal site. (c) There is no analysis on the quality of the reported optimal road segment, either theoretically or empirically. (d) It is not shown how does the reported road segment performs for other established metrics, such as number of new users covered, distance traveled by the users to avail the service, etc.

Facility location problems [13, 18] typically consider a set of users, and a set of candidate sites. The goal is to identify a set of k candidate sites that optimize certain metrics such as covering maximum number of users, or minimizing the average distance between a user and its nearest facility, etc. Almost all of these problems are NP-hard. While early works assumed that the users are static, mobile users are now considered. The *flow based* facility location works [2–7] assume a flow model to characterize human mobility, instead of using real trajectories. A fairly comprehensive literature survey is available in [8]. We briefly outline the major works related to the different versions of TOPS queries that have been discussed in Sec. 7. In [7], few exact and approximate algorithms for TOPS1 and TOPS4, were presented, under the restriction that a customer would stay on the path, i.e., $\tau = 0$. Later, in [2], few generalizations of the model were proposed, where the customers were allowed to deviate. These include TOPS1, TOPS2 and TOPS3 problems. Further generalizations of TOPS1 were examined in [6], with probabilistic customer flows, single and multiple flow interceptions, and fixed and variable installation costs of the services. Several existing flow-based facility location models were generalized in [43]. Few flow refueling location models have been proposed in [26,31,34] for sitingalternative- fuel stations. We have already discussed how our work differs from these works in Sec. 1.

Recently, a greedy heuristic was proposed for the TOPS1 problem in [29]. under the assumption that $\tau = 0$, i.e., no customer deviations are allowed. We present a detailed study of the TOPS1 problem in [36], with arbitrary value of τ.

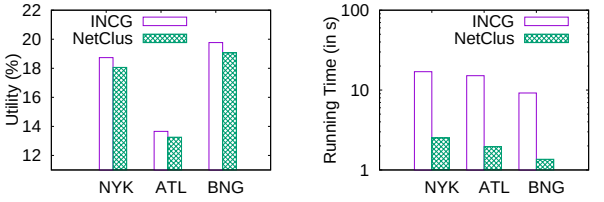

(a) Utility. (b) Running time.

Figure 11: Effect of city geometries ($k = 5$ and $\tau = 0.8$ Km).

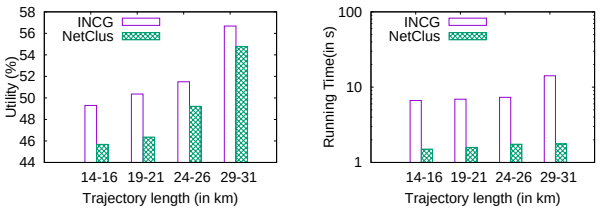

(a) Utility. (b) Running time.

Figure 12: Effect of trajectory length ($k = 5$ and $\tau = 0.8$ Km).

R_p (Km)	κ	$\lvert\bar{\Lambda}\rvert$	$\lvert\bar{\mathcal{TL}}\rvert$	$\lvert\bar{\mathcal{CL}}\rvert$	Run-time (s)
0.0093	258340	1.04	561.88	4.29	269.35
0.0286	195910	1.15	571.41	6.43	239.58
0.0163	233729	1.38	592.45	12.62	255.60
0.0500	153210	1.76	626.18	19.44	221.79
0.0875	112223	2.40	675.60	26.04	204.38
0.1531	76836	3.51	757.03	45.38	192.12
0.2680	48288	5.58	895.19	64.27	188.25
0.4689	28510	9.46	1162.93	53.74	205.74
0.8207	15775	17.10	1525.09	42.73	281.62
1.4361	8258	32.66	2162.01	35.62	537.27
2.5133	4202	64.18	3092.64	21.73	1300.12
4.3982	2024	133.24	4148.03	12.49	3231.92
7.6968	938	287.51	7537.33	3.51	7333.60

Table 11: Details of indexing for Beijing road network comprising of 269,686 nodes with $\gamma = 0.75$.

10. CONCLUSIONS

In this paper, we have proposed a generic TOPS framework to solve the problem of finding facility locations for trajectory-aware services. We showed that the problem is NP-hard, and proposed a greedy heuristic with constant factor approximation bound. However, this heuristic does not scale for large datasets. Thus, we developed an index structure, NETCLUS, to make it practical for city-scale road networks. Extensive experiments over real datasets showed that NETCLUS yields solutions that are comparable with those of the greedy heuristic, while being significantly faster, and low in memory overhead. The proposed framework can handle a wide class of objectives, and additional constraints, thus making it highly generic and practical.

11. REFERENCES

[1] P. Banerjee, S. Ranu, and S. Raghavan. Inferring uncertain trajectories from partial observations. In *ICDM*, pages 30–39, 2014.

[2] O. Berman, D. Bertsimas, and R. C. Larson. Locating discretionary service facilities, ii: maximizing market size, minimizing inconvenience. *Operations Research*, 43(4):623–632, 1995.

[3] O. Berman and D. Krass. Flow intercepting spatial interaction model: a new approach to optimal location of competitive facilities. *Location Science*, 6(1):41–65, 1998.

[4] O. Berman and D. Krass. The generalized maximal covering location problem. *Computers & Operations Research*, 29(6):563–581, 2002.

[5] O. Berman, D. Krass, and C. W. Xu. Locating discretionary service facilities based on probabilistic customer flows. *Transportation Science*, 29(3):276–290, 1995.

[6] O. Berman, D. Krass, and C. W. Xu. Locating flow-intercepting facilities: New approaches and results. *Annals of Operations research*, 60(1):121–143, 1995.

[7] O. Berman, R. C. Larson, and N. Fouska. Optimal location of discretionary service facilities. *Transportation Science*, 26(3):201–211, 1992.

[8] M. Boccia, A. Sforza, and C. Sterle. Flow intercepting facility location: Problems, models and heuristics. *J. Mathematical Modelling and Algorithms*, 8(1):35–79, 2009.

[9] L. Chen and R. Ng. On the marriage of edit distance and Lp norms. In *VLDB*, 2004.

[10] N. Chen, J. Meng, J. Rong, and H. Zhu. Approximation for dominating set problem with measure functions. *Computing and Informatics*, 23(1):37–49, 2012.

[11] Z. Chen, Y. Liu, R. C.-W. Wong, J. Xiong, G. Mai, and C. Long. Efficient algorithms for optimal location queries in road networks. In *SIGMOD*, pages 123–134, 2014.

[12] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2009.

[13] Z. Drezner. *Facility location: a survey of applications and methods*. Springer, 1995.

[14] Y. Du, D. Zhang, and T. Xia. The optimal-location query. In *Advances in Spatial and Temporal Databases*, pages 163–180, 2005.

[15] P. Flajolet and G. N. Martin. Probabilistic counting algorithms for data base applications. *J. Computer and System Sciences*, 31(2):182–209, 1985.

[16] E. Frentzos, K. Gratsias, and Y. Theodoridis. Index-based Most Similar Trajectory Search. In *ICDE*, pages 816–825, 2007.

[17] P. Ghaemi, K. Shahabi, J. P. Wilson, and F. Banaei-Kashani. Optimal network location queries. In *SIGSPATIAL GIS*, pages 478–481, 2010.

[18] H. W. Hamacher and Z. Drezner. *Facility location: applications and theory*. Springer, 2002.

[19] T. W. Haynes, S. Hedetniemi, and P. Slater. *Fundamentals of Domination in Graphs*. CRC Press, 1998.

[20] W. He, D. Li, T. Zhang, L. An, M. Guo, and G. Chen. Mining regular routes from gps data for ridesharing recommendations. In *SIGKDD*, pages 79–86, 2012.

[21] H. Jeung, M. L. Yiu, X. Zhou, C. S. Jensen, and H. T. Shen. Discovery of convoys in trajectory databases. *Proceedings of the VLDB Endowment*, 1(1):1068–1080, 2008.

[22] P. Kalnis, N. Mamoulis, and S. Bakiras. On discovering moving clusters in spatio-temporal data. In *International Symposium on Spatial and Temporal Databases*, pages 364–381. Springer, 2005.

[23] B. kee Yi, H. V. Jagadish, and C. Faloutsos. Efficient Retrieval of Similar Time Sequences Under Time Warping. In *ICDE*, 1998.

[24] S. Khuller, A. Moss, and J. S. Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.

[25] V. Kolar, S. Ranu, A. P. Subramanian, Y. Shrinivasan, A. Telang, R. Kokku, and S. Raghavan. People in motion: Spatio-temporal analytics on call detail records. In *COMSNETS*, pages 1–4, 2014.

[26] M. Kuby, L. Lines, R. Schultz, Z. Xie, J.-G. Kim, and S. Lim. Optimization of hydrogen stations in florida using the flow-refueling location model. *International journal of hydrogen energy*, 34(15):6045–6064, 2009.

[27] J.-G. Lee, J. Han, and K.-Y. Whang. Trajectory clustering: a partition-and-group framework. In *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*, pages 593–604. ACM, 2007.

[28] X. Li, V. Čeikute, C. S. Jensen, and K.-L. Tan. Trajectory based optimal segment computation in road network databases. In *SIGSPATIAL GIS*, pages 396–399, 2013.

[29] Y. Li, J. Bao, Y. Li, Y. Wu, Z. Gong, and Y. Zheng. Mining the most influential k-location set from massive trajectories. In *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, page 51. ACM, 2016.

[30] Z. Li, B. Ding, J. Han, and R. Kays. Swarm: Mining relaxed temporal moving object clusters. *Proceedings of the VLDB Endowment*, 3(1-2):723–734, 2010.

[31] S. Lim and M. Kuby. Heuristic algorithms for siting alternative-fuel stations using the flow-refueling location model. *European Journal of Operational Research*, 204(1):51–61, 2010.

[32] X. Lin, Y. Yuan, Q. Zhang, and Y. Zhang. Selecting stars: The k most representative skyline operator. In *ICDE*, pages 86–95, 2007.

[33] Y. Lou, C. Zhang, Y. Zheng, X. Xie, W. Wang, and Y. Huang. Map-matching for low-sampling-rate gps trajectories. In *GIS*, pages 352–361, 2009.

[34] S. MirHassani and R. Ebrazi. A flexible reformulation of the refueling station location problem. *Transportation Science*, 47(4):617–628, 2012.

[35] S. Mitra, S. Ranu, V. Kolar, A. Telang, A. Bhattacharya, R. Kokku, and S. Raghavan. Trajectory aware macro-cell planning for mobile users. In *INFOCOM*, 2015.

[36] S. Mitra, P. Saraf, R. Sharma, A. Bhattacharya, S. Ranu, and H. Bhandari. NetClus: A scalable framework for locating top-k sites for placement of trajectory-aware services. *ICDE*, 2017.

[37] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, 1978.

[38] S. Ranu, P. Deepak, A. D. Telang, P. Deshpande, and S. Raghavan. Indexing and matching trajectories under inconsistent sampling rates. In *ICDE*, 2015.

[39] M. Vlachos, D. Gunopoulos, and G. Kollios. Discovering similar multidimensional trajectories. In *ICDE*, 2002.

[40] X. Xiao, B. Yao, and F. Li. Optimal location queries in road network databases. In *ICDE*, pages 804–815, 2011.

[41] J. Yuan, Y. Zheng, X. Xie, and G. Sun. Driving with knowledge from the physical world. In *KDD*, pages 316–324, 2011.

[42] J. Yuan, Y. Zheng, C. Zhang, W. Xie, X. Xie, G. Sun, and Y. Huang. T-drive: driving directions based on taxi trajectories. In *SIGSPATIAL GIS*, pages 99–108, 2010.

[43] W. Zeng, I. Castillo, and M. J. Hodgson. A generalized model for locating facilities on a network with flow-based demand. *Networks and Spatial Economics*, 10(4):579–611, 2010.

[44] W. Zeng, M. J. Hodgson, and I. Castillo. The pickup problem: consumers' locational preferences in flow interception. *Geographical Analysis*, 41(1):107–126, 2009.

[45] D. Zhang, Y. Du, T. Xia, and Y. Tao. Progressive computation of the min-dist optimal-location query. In *Proceedings of the 32nd international conference on Very large data bases*, pages 643–654. VLDB Endowment, 2006.

τ (Km)	Running Time (s)	Memory (in GB)
0	2890	3.88
0.2	12473	8.93
0.4	26693	9.98
0.8	28391	14.23
1.2	29058	16.57
1.6	32647	20.85
2.0	34618	26.74
2.4	Out of memory	

Table 12: Performance of Jaccard similarity based clustering for Beijing road network using a Jaccard distance threshold $\alpha = 0.8$.

Algorithm 1 INC-GREEDY

```
1: procedure INC-GREEDY(k, τ, ψ)
2:     Compute the sets TC, SC and the site-weights.
3:     Q_0 ← ∅
4:     for each s_i ∈ S do
5:         U_0(s_i) ← w_i
6:         ∀T_j ∈ TC(s_i), α_ji ← ψ(T_j, s_i)
7:     ∀T_j ∈ T, U_j ← 0
8:     for each θ = 1, ..., k do
9:         s_θ ← arg max_{∀s_i ∈ S \ Q_{θ-1}} {U(Q_{θ-1} ∪ {s_i}) − U(Q_{θ-1})}
10:        Q_θ ← Q_{θ-1} ∪ {s_θ}
11:        for each trajectory T_j ∈ TC(s_θ) do
12:            U_j ← max(U_j, ψ(T_j, s_θ))
13:            if U_j changes then
14:                for each site s_i ∈ SC(T_j), s_i ∉ Q_θ do
15:                    α'_ji ← max(0, ψ(T_j, s_i) − U_j)
16:                    U_θ(s_i) ← U_{θ-1}(s_i) − (α_ji − α'_ji)
17:                    α_ji ← α'_ji
```

APPENDIX

A. ALGORITHMS FOR TOPS

A.1 Linearization for Optimal Algorithm

Each inequality in Ineq. (6) can be expressed as

$$U_j \leq \max\{U_{j1}, U_{j2}\}$$

where

$$U_{j1} \leq \max\{\psi(T_j, s_i)x_i | 1 \leq i \leq \lfloor n/2 \rfloor\}$$

and

$$U_{j2} \leq \max\{\psi(T_j, s_i).x_i | \lfloor n/2 \rfloor + 1 \leq i \leq n\}$$

The terms U_{j1} and U_{j2} can be recursively expressed in the same manner as in case of U_j.

Finally, the constraint $U_j \leq \max\{U_{j1}, U_{j2}\}$ can be linearized as follows:

$$\begin{aligned} U_{j1} &\leq U_{j2} + M.y_j \\ U_{j2} &\leq U_{j1} + M.(1 - y_j) \\ U_j &\leq U_{j2} + M.y_j \\ U_j &\leq U_{j1} + M.(1 - y_j) \\ y_j &\in \{0, 1\} \end{aligned}$$

where M is a sufficiently large number.

A.2 Inc-Greedy Algorithm

The INC-GREEDY procedure is outlined in Algorithm 1. The algorithm begins by computing the sets TC, SC, and the site-weights. It then initializes the marginal utilities of each site $s_i \in \mathcal{S}$ to its site weight, i.e., $U_0(s_i) = w_i$. Further, for each trajectory $T_j \in TC(s_i)$, it initializes $\alpha_{ji} = \psi(T_j, s_i)$.

In iteration $\theta = 1, \ldots, k$, it selects the site s_θ with *maximal* marginal utility, i.e., $\forall s_i \notin \mathcal{Q}_{\theta-1}, U_\theta(s_\theta) \geq U_\theta(s_i)$. If multiple candidate sites yield the same maximal marginal gain we choose the one with maximal weight. Still, if there are ties, then without loss of generality, the site with the highest index is selected.

For each trajectory $T_j \in TC(s_\theta)$ (line 11), first the utility U_j is updated (line 12). In case there is a gain in U_j, each site $s_i \in SC(T_j)$ is processed as indicated in line 14. In line 15, the new marginal utility of trajectory T_j w.r.t. site s_i is computed. Using this value, the marginal utility of site s_i and α_{ji} are updated.

B. NETCLUS

B.1 Clustering based on Jaccard Similarity

Given two nodes $s_i, s_j \in \mathcal{S}$, the Jaccard similarity between their trajectory covers is

$$J_s(s_i, s_j) = \frac{|TC(s_i) \cap TC(s_j)|}{|TC(s_i) \cup TC(s_j)|}$$

The Jaccard distance is

$$J_d(s_i, s_j) = 1 - J_s(s_i, s_j)$$

The nodes in the road network V are clustered using the above measure as follows.

The node v_1 with the highest weight, i.e., the one whose sum of preference scores over the covered trajectories as defined in Sec. 3.2, is chosen as the first cluster center of the cluster g_1. All nodes that are within a threshold of α Jaccard distance from v_1 are added to g_1. From the remaining set of unclustered nodes, the node with the highest weight is chosen as a cluster center, and the process is repeated until each node is clustered.

The results of clustering on the Beijing dataset are shown in Table 12. The limitations of Jaccard similarity based clustering are explained in Sec. 4.

B.2 Complexity Analysis of Greedy-GDSP

THEOREM 10. *Greedy-GDSP runs in $O(|V|.(\nu \log \nu + \eta))$ time where ν is the maximum number of vertices that are reachable within the largest round-trip distance R_{max} from any vertex v, and η is the number of clusters returned by the algorithm.*

PROOF. As the underlying graph models a road network which is roughly planar, we assume that the number of edges (i.e., road segments) incident on any set of ν vertices is $O(\nu)$. Since the cost of running the shortest path algorithm from a given node in a graph $G = (V, E)$ is $O(|E| \log |V|)$ [12], therefore, the initial distance computation for a given source vertex requires $O(\nu \log \nu)$ time. Thus, the distance computation across all the vertices in V takes $O(|V|.\nu \log \nu)$. The neighbors of each node v can be maintained as a list sorted by the round-trip distance from v. Therefore, computing the dominating set for a particular R requires at most ν time for each vertex. The total time for the construction of dominating sets, hence, is $O(|V|.\nu)$.

Choosing the vertex with the largest *incremental* dominating set requires bitwise OR operations of FM sketches. For $|V|$ vertices, there are at most $|V|-1$ such operations, each requiring $O(f)$ time (since there are f copies of FM sketches each with a constant size of 32 bits). As the number of clusters produced is η, the running time is $O(|V|.\eta)$.

The total running time is, therefore, $O(|V|.(\nu \log \nu + \eta))$. □

B.3 Complexity of NetClus

THEOREM 11. *The time and space complexities of* NETCLUS *are $O(k.\eta_p.\xi_p)$ and $O(\sum_{p=1}^{t}(\eta_p(\xi_p + \lambda_p)))$ respectively.*

PROOF. The number of cluster representatives, $|\widehat{\mathcal{S}}|$ is at most the number of clusters, η_p. For any cluster $g_i \in \mathcal{I}_p$, $|\mathcal{TL}(g_i)| =$

$O(\xi_p)$. We assume that the number of neighboring clusters for any cluster is bounded by a constant, i.e., $|\mathcal{CL}(g_i)| = O(1)$. (Table 11 that lists typical mean values of $|\mathcal{CL}(g_i)|$ empirically supports the assumption.) Hence, computing the set $\widehat{TC}(r_i)$ requires at most $O(\xi_p)$ time. The total time across all the η_p clusters, therefore, is $O(\eta_p.\xi_p)$. The inverse maps $\widehat{CC}(T_j)$ for the trajectories can be computed along with $\widehat{TC}(r_i)$. Hence, the total time remains $O(\eta_p.\xi_p)$. Using Th. 4, the subsequent INC-GREEDY phase requires $O(k.\eta_p.\xi_p)$ time. Therefore, the overall time complexity of the algorithm is $O(k.\eta_p.\xi_p)$.

We next analyze the space complexity for a particular index instance $\mathcal{I}_p$. Storing the sets $\mathcal{TL}(g_i)$, each of size at most ξ_p, across all the η_p clusters require $O(\eta_p.\xi_p)$ space. Storing the inverse maps $CC(T_j)$ for all the trajectories, thus, also requires $O(\eta_p.\xi_p)$ space. As there are at most λ_p vertices in any cluster, storing their ids and distances from the cluster center requires $O(\lambda_p)$ space for each cluster. Assuming $|\mathcal{CL}(g_i)| = O(1)$, the total storage cost for a cluster is $O(\lambda_p)$. Therefore, the total storage cost is $O(\eta_p.(\xi_p + \lambda_p))$. Summing across all the index instances, the entire space complexity is $O(\sum_{p=1}^{t}(\eta_p(\xi_p + \lambda_p)))$. □